United States
Department
of Agriculture

Forest Service

Intermountain
Research Station

Research Paper
INT-447

July 1991

A Heuristic Process for Solving Mixed-Integer Land Management and Transportation Planning Models

J. Greg Jones
Andres Weintraub
Mary L. Meacham
Adrian Magendzo

THE AUTHORS

J. GREG JONES is a research forester with the Economics Research Work Unit, Forestry Sciences Laboratory, Missoula, MT. His research includes economic efficiency modeling for use in multiple-use management and the measurement of multiple-use costs and benefits. He received his Ph.D. degree in forest economics from Iowa State University in 1976.

ANDRES WEINTRAUB is Professor, Department of Industrial Engineering, University of Chile. His research includes forest plan models and algorithms, operations management, and mathematical programming. He received his Ph.D. degree in industrial engineering and operations research from the University of California, Berkeley in 1971.

MARY L. MEACHAM is a statistician with the Economics Research Work Unit, Forestry Sciences Laboratory, Missoula, MT. Her research includes work in mathematical programming, regression analysis, and other statistical techniques. She received her B.A. degree in mathematics from the University of Montana in 1985.

ADRIAN MAGENDZO is research associate, Department of Industrial Engineering, University of Chile. His research includes work in system design and implementation of operations research models. He received his industrial engineering degree from the University of Chile in 1986.

RESEARCH SUMMARY

The mathematical programming formulation first developed by the Integrated Resource Planning Model (IRPM) (Kirby and others 1980) is useful in identifying the spatial arrangement and timing of land management activities and road construction and reconstruction projects that efficiently implement management objectives. In addition to IRPM, currently this formulation can also be developed using the Integrated Resource Analysis System (IRAS) (Jones and others 1990), or the project option within FORPLAN (Johnson and others 1986).

The IRPM-type formulation is a mixed-integer mathematical program; that is, both continuous and 0,1 integer variables are present. The integer variables arise in the transportation portion of the formulation. Road construction and reconstruction projects are developed for specific segments of proposed and existing roads. These decision variables are formulated as integers that can only assume the values of 1 (road project is selected) or 0 (road project is not selected). Fractional values for road projects are not allowed because they represent the construction of only portions of road segments. This leads to a discontinuous road network, which is an impractical solution.

Unfortunately, mixed-integer mathematical programming problems are, in general, difficult to solve, and this has hampered use of the IRPM-type formulation. Experience has shown that branch-and-bound algorithms for solving these mixed-integer formulations, such as the algorithm available at National Computer Center at Fort Collins (NCC-FC), are viable for solving only relatively small models. For models typical in real-world applications containing 200 or more integer variables, the branch-and-bound approach is generally cost-prohibitive, if optimal solutions can be obtained at all.

This paper describes a heuristic procedure for solving the IRPM-type formulation. The procedure is an iterative process in which models are initially formulated and solved as continuous linear programming problems (road projects are represented by continuous variables). Then logical decision rules (based on the structure of the IRPM-type formulation) are applied to the continuous linear programming (LP) solution to round some of the fractional road projects to 0 or 1. The objective in these rounding decisions is to preserve feasibility and obtain solutions as close as possible to optimum. These rounding decisions are incorporated into the LP matrix and another LP solution is made. Based on this new solution, rounding decisions made in previous iterations are reviewed to determine if any changes should be made and, secondly, additional fractional projects are rounded to 0 or 1. These decisions are incorporated into the LP matrix, another LP solution is obtained, and so on. The process continues until no fractional road projects remain in the LP solution and no changes in the previous decisions can be found that would improve the value of the objective function. Feasible mixed-integer solutions for full-scale models can generally be obtained within seven to 12 iterations.

Intermountain Research Station
324 25th Street
Ogden, UT 84401

This heuristic procedure has been coded in FORTRAN and is part of a set of software routines that reside at NCC-FC. The procedure is automated in the sense that all iterations are handled without user intervention. Full-scale models containing up to 4,000 rows, an equal number of columns, and 300 or more integer variables can be solved with about 20 minutes of central processing unit (CPU) time, with costs ranging from $50 to $200, given July 1990 rates at NCC-FC. IRPM-type formulations developed with IRPM, IRAS, and FORPLAN can be solved with this procedure.

The advantage of this heuristic solution procedure is that feasible mixed-integer solutions for the IRPM-type formulation can be obtained for full-scale models with significantly less computer time than what is required to solve the same problem using the branch-and-bound approach. But because the solution procedure is heuristic, it does not provide optimal mixed-integer solutions. Comparisons, however, with the optimal mixed-integer solutions and the objective value of first continuous solutions made in the heuristic process (which is an upper bound for the objective value for the optimal mixed-integer solution) indicate that the objective value of heuristic solutions is generally well within 10 percent of the objective value of the optimal mixed-integer solution. Other comparisons have been made with a currently used planning method in which managers choose land management activities and use a road network model for finding the least-cost access routes. In these comparisons, the objective values of the plans developed via the heuristic solution procedure exceeded the other planning approach by an average of 40 percent, while satisfying the same stated management objectives.

ACKNOWLEDGMENT

This study was partly funded by FONDECYT Chile under Grant 1070-90.

CONTENTS

Page

Introduction 1
The IRPM-Type Formulation 2
  Decision Variables 5
  Constraints 5
  Objective Function 7
The Heuristic Integer Programming (HIP) Procedure 8
  Formulation Modifications 8
  Iteration 1: Actions Following the First LP Solution 9
  Iteration 2: Actions Following the Second LP Solution 11
  Iteration 3: Actions Following the Third LP Solution 12
  Iteration 4: Actions Following the Fourth LP Solution 14
  Iteration 5: Actions Following the Fifth LP Solution 15
  Iterations 6+: Actions following the Sixth to the Last LP Solution 16
Solution Validation 16
  Comparison With Optimal Mixed-Integer Solutions 16
  Comparison With Other Planning Methods 20
  Comparison With the First Continuous LP Solution 21
  Computer Resource Requirements 22
Conclusions 23
References 24
Appendix A: Capacity Adjustment #1 25
Appendix B: Capacity Adjustment #2 26
Appendix C: Rounding Procedure 27
  Develop a List of Candidate Road Projects for Rounding 28
  Calculate Index Values for the Candidate Projects 30
  "Tree-Growth" Rounding Procedure 32
Appendix D: Feasibility Checking 35
  Comparing Against "Actual Space" 35
  Comparing Against "Allotted Space" 37

A Heuristic Process for Solving Mixed-Integer Land Management and Transportation Planning Models

J. Greg Jones
Andres Wientraub
Mary L. Meacham
Adrian Magendzo

INTRODUCTION

Linear programming applications for timber harvest scheduling and resource allocation have been common in forestry since the early 1970's. Examples include FORPLAN (Johnson and others 1986) and earlier models such as MUSYC (Johnson and Jones 1979), Timber RAM (Navon 1971), and the Resource Allocation Analysis system (USDA Forest Service 1975). These applications have been oriented towards strategic planning where time frames are quite long and the areas analyzed quite large. They have been used to develop the standards and guidelines by which forest land is to be managed, as well as output targets. They do not, however, identify where or when to conduct projects to achieve these objectives.

Tactical planning is concerned with the arrangement and timing of projects that best meet the management direction and objectives of a strategic plan. There are often numerous considerations in making tactical timber harvest and transportation decisions for implementing long-range strategic plans. These decisions address the location of the harvest units, silvicultural methods, logging methods, and when and how to best access the harvest units. Decisions about such things affect the profitability of a timber sale, the environmental effects, and the extent to which multiple-use objectives are met. But it does not end there. Decisions about a sale also have implications regarding future management within the general area surrounding the sale. The geographic placement of the harvest units in the current sale area will affect where harvest units can be located in the future because of impacts on critical watersheds and other cumulative environmental effects, wildlife dispersion objectives, and concerns about esthetics. Current activities may also affect the financial aspects of selling the residual timber in the future.

The Integrated Resource Planning Model (IRPM) (Kirby 1980), the Integrated Resource Analysis System (IRAS) (Jones and others 1990), and the project option within FORPLAN (Johnson and others 1986) can be used to develop a formulation designed for tactical forest planning. Tests have indicated that this formulation is quite useful in identifying land management and transportation projects that best meet specified management objectives (Jones and others 1986).

One problem that has hampered widespread use of the IRPM-type formulation is the lack of an efficient, easy-to-use solution procedure. Some of the

decision variables in this formulation are integers that can assume only the values of 0 or 1. Experience has shown that the branch-and-bound approach for solving mixed-integer formulations (such as the Functional Mathematical Programming System [Sperry Corporation 1984], the mathematical programming system at the USDA National Computer Center at Fort Collins) are viable for solving only relatively small models. For models typical in real-world applications containing 200 or more integer variables, the branch-and-bound approach has not been practical.

This paper describes a heuristic procedure for solving the IRPM-type mixed-integer formulation. It is based on an approach originally proposed by Weintraub (1982). The procedure has been coded in a program that operates at the USDA National Computer Center at Fort Collins (NCC-FC). This heuristic solution procedure can process models developed with IRPM, IRAS, or FORPLAN using the project analysis formulation. (Readers interested in more information about this computer program should contact the authors.)

The advantage afforded by this heuristic solution procedure is that feasible mixed-integer solutions can be found for large, real-world problems with much less computer time than required by a branch-and-bound algorithm. Models in the range of 2,000 to 4,000 rows, an equal number of columns, and 200 or more integer variables can usually be solved within 20 minutes of central processing unit (CPU) time on the UNISYS mainframe computer at NCC-FC. Costs range from $50 to $200, given the current rate structure at NCC-FC.

The heuristic procedure does not obtain optimal mixed-integer solutions. Tests, however, comparing the heuristic procedure with a branch-and-bound algorithm and with other methods for tactical planning have shown that it provides good mixed-integer solutions, often within 10 percent of the objective function value of the true mixed-integer optimum.

The paper begins with a presentation of the IRPM-type formulation this heuristic process is designed to solve. Next, the heuristic process is presented, but largely in descriptive terms. Much of the detail regarding computations is presented in appendices A-D. Last, results of the heuristic process are compared with results of a branch-and-bound algorithm and with other planning methods.

THE IRPM-TYPE FORMULATION

The IRPM-type formulation combines a land allocation and scheduling model with a cost-minimizing transportation network model. Land is delineated into contiguous tracts, hereafter called polygons. These polygons may represent potential harvest units (as illustrated in fig. 1), timber stands, portions of stands, or other land delineations that result in contiguous tracts. Land management projects, called "resource projects," are then formulated for the polygons. These projects vary by the type and timing of management treatments to be applied. To access these polygons, a road network is developed as illustrated in figure 2. This network is divided into "links," which are bounded on the ends by "nodes." Road construction projects are specified for links representing proposed new roads, and reconstruction projects are specified for links representing existing roads requiring improvement prior to some specified future use. These projects vary by the standard to which the road would be constructed and the timing for

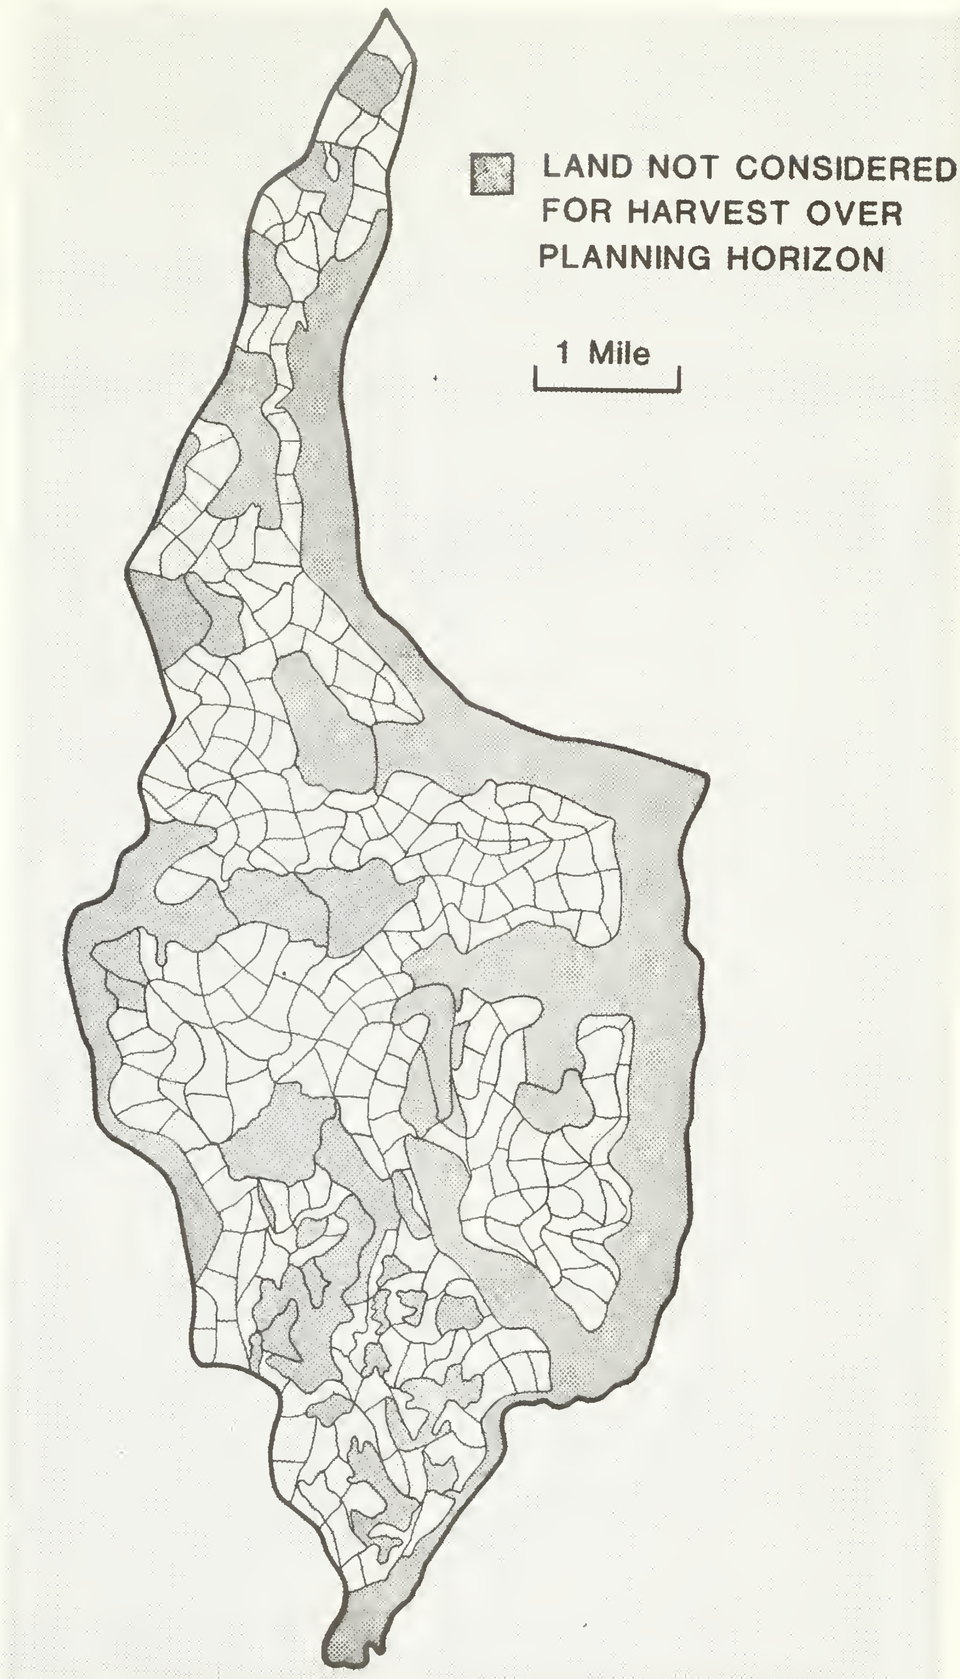

Figure 1—Example of polygons delineated for an area.

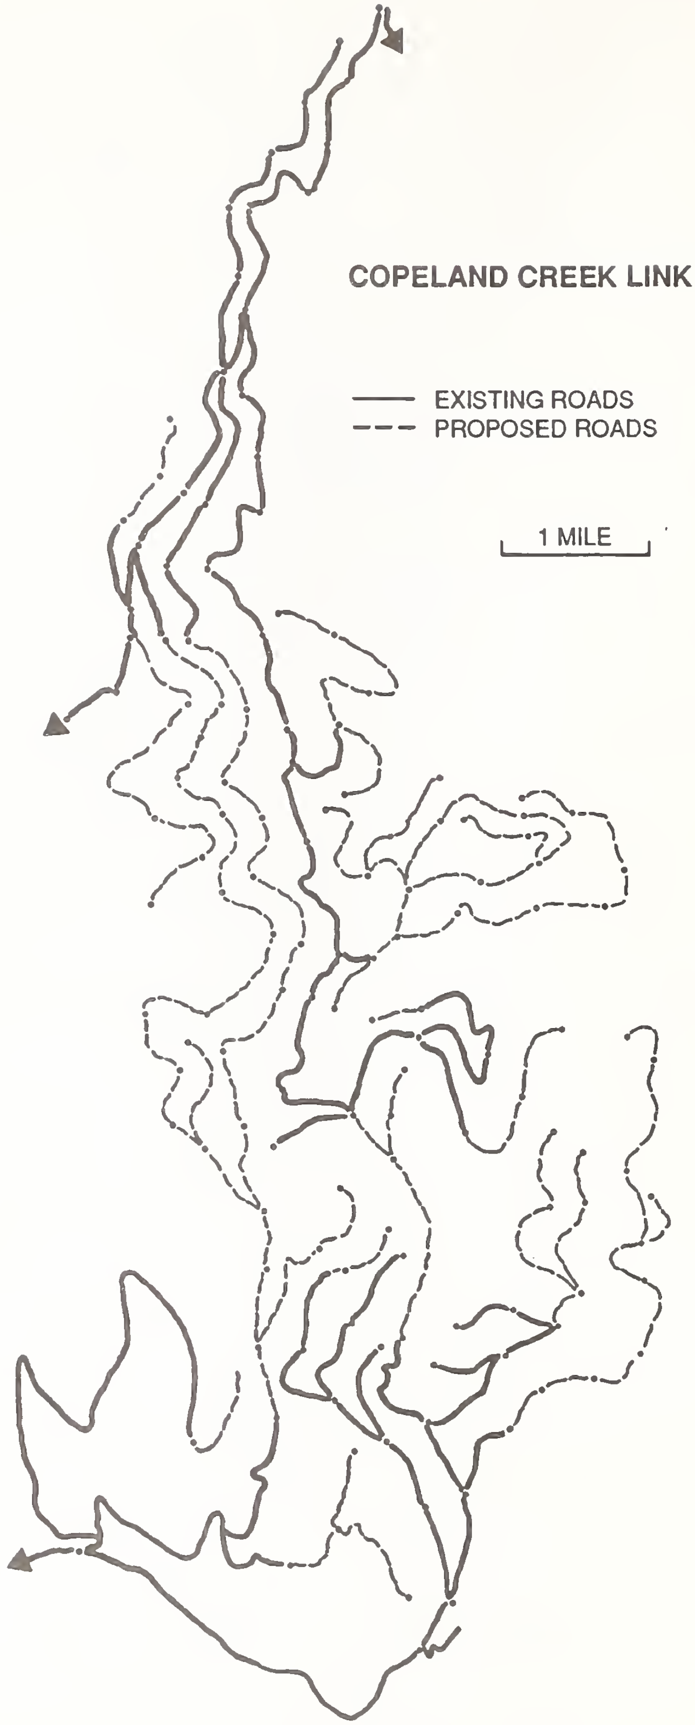

Figure 2—Example of a proposed road network for an area.

construction. Traffic in the model is assumed to originate at the site of land management activities, and is channeled through the network to one or more "final demand" nodes on the extremities of the network. Round-trip impacts and costs are achieved by assuming traffic returns via the same route. The basic structure of the model is presented below.

Decision Variables

1. Resource Projects:

X_{lit} = the fraction of the resource project comprised of management alternative l to be implemented on polygon i in period t. ($0 \leq X_{lit} \leq 1$)

2. Road Projects:

$$Z^{k}_{ab,t} = \begin{cases} 1 & \text{if the road link connecting adjacent nodes } a \text{ and } b \text{ is to be constructed (or reconstructed) to standard } k \text{ in period } t. \text{ For existing roads, a value of 1 means standard } k \text{ (which has no construction cost) is available for use starting in period } t. \\ 0 & \text{otherwise} \end{cases}$$

3. Traffic Flow Variables:

$T^{k,y}_{ab,t}$ = amount of traffic type y (for example, logging truck or recreation) on road standard k flowing from node a to node b in period t. (A return trip in the opposite direction is implied.)

$T^{k,y}_{bc,t}$ = amount of traffic type y (for example, logging truck or recreation) on road standard k flowing from node b to node c in period t. (A return trip in the opposite direction is implied.)

$D^{y}_{bf,t}$ = amount of traffic type y flowing from node b and terminating at final demand node f in period t. (This is, in actuality, a special type of traffic variable that denotes a terminating point in the road network.)

Constraints

1. Mutually exclusive constraints for road construction projects:

$$\sum_{k} \sum_{t} Z^{k}_{ab,t} \leq 1 \qquad \text{(for all } ab\text{)} \quad (1)$$

Only one road project may be selected for any road link, ab.

2. Traffic flow equations:

Each traffic flow equation equates the volume of traffic type y in period t entering node b to the volume of traffic type y leaving that node in period t. In doing so the traffic generated by the resource projects is routed through the road network to the final demand locations. They are formulated as:

$$\sum_{l} \sum_{i} V^{y}_{b,lit}(X_{lit}) + \sum_{a \varepsilon A} \sum_{k} (T^{k,y}_{ab,t})$$

$$= \sum_{c \varepsilon C} \sum_{k} (T^{k,y}_{bc,t}) + \sum_{f \varepsilon F} (D^{y}_{bf,t}) \qquad \text{(for all } b,y,t\text{)} \quad (2)$$

where

$V^{y}_{b,lit}$ = amount of traffic type y loaded by resource project X_{lit} onto node b in period t ($V^{y}_{b,lit}$ is zero for resource projects that do not load traffic onto node b)

A = the set of all nodes a from which traffic may travel (over one road link) to node b

C = the set of all nodes c to which traffic may travel (over one road link) from node b

F = the set of all final demand nodes f connected to node b by a road link.

3. Capacity equations:

$$\sum_{y} [E_y (T^{k,y}_{ab,t}) + E_y (T^{k,y}_{ba,t})] \leq \sum_{g=1}^{t} [C^{k}_{ab,g} (Z^{k}_{ab,g})] \qquad \text{(for all } k,ab,t\text{)} \quad (3)$$

where

E_y = amount one unit of traffic type y counts toward capacity on link ab

$C^{k}_{ab,g}$ = capacity of standard k (measured in a common unit of traffic) on link ab in time g.

Capacity constraints serve two purposes. First, if link ab is to carry traffic in period t, the capacity equations ensure one road project is selected for link ab by period t (that is, one of the road construction projects in one of the capacity equations for link ab must assume a value of 1). Second, because each $Z^{k}_{ab,g}$ is restricted to the values of 0 or 1 and the sum of these variables defined for link ab cannot exceed 1, capacity constraints limit the amount of traffic allowed in each period g for road standard k on link ab to $C^{k}_{ab,g}$.

4. Mutually exclusive constraints for resource projects:

$$\sum_{l} \sum_{t} X_{lit} \leq 1 \qquad \text{(for all } i\text{)} \quad (4)$$

Only one or the equivalent of one resource project may be selected for any polygon, i.

5. Typical optional relationships among resource projects:

(i) Mutually exclusive relationships specify that the equivalent of one project may be selected from a user-defined set S:

$$\sum_{(l,i,t)\varepsilon S} X_{lit} \leq 1 \qquad (5)$$

(ii) Contingent relationships require that at least one project from a user-defined set S must be selected if the user-specified project X_{lit}* is selected:

$$X_{lit}^* \leq \sum_{(l,i,t)\varepsilon S} X_{lit} \qquad (6)$$

(iii) Companion relationships require that if user-specified project X_{lit} is selected, then some second user-specified project X_{lit}* must also be selected at the same level, and vice versa:

$$X_{lit} = X_{lit}^* \qquad (7)$$

6. Resource Project to Road Construction Triggers:

$$X_{lit} \leq \sum_{(ab,k)\varepsilon S} \sum_{g=1}^{t} Z_{ab,g}^{k} \tag{8}$$

where S is comprised of a set of road projects, one of which must be selected if X_{lit} is selected. These triggers are used to ensure one or more specific road segments in the vicinity of X_{lit} are constructed if X_{lit} is selected. They are not used to specify routes to a final demand node, however. This is accomplished via the traffic flow equations (equation 2) and capacity equations (equation 3) which, respectively, (1) channel all traffic originating with the selected resource projects to one or more final demand nodes, and (2) require that a road project be selected for any link to carry traffic.

7. Side Constraints:

The term "side constraints" refers to a class of user-defined management constraints. Side constraints include constraints placed on costs (such as road construction costs, hauling costs, or total discounted costs), on physical outputs (such as timber yields, water production, or sediment produced), and on net revenues. These constraints can contain any of the decision variable types or any combination of them and can be formulated as upper limit, lower limit, or equality constraints.

Objective Function

Various objectives may be defined. A common one is to maximize present net value (PNV):

$$\begin{aligned} \max \text{PNV} = & \sum_{l} \sum_{i} \sum_{t} G_{lit}(X_{lit}) - \sum_{k} \sum_{ab} \sum_{t} H_{ab,t}^{k}(Z_{ab,t}^{k}) \\ & - \sum_{k} \sum_{y} \sum_{ab} \sum_{t} I_{ab,t}^{k,y}(T_{ab,t}^{k,y}) \\ & - \sum_{k} \sum_{y} \sum_{ba} \sum_{t} I_{ba,t}^{k,y}(T_{ba,t}^{k,y}) \end{aligned} \tag{9}$$

where

G_{lit} = net value of management alternative l on polygon i in period t. For timber, G_{lit} is computed as the expected value of the delivered logs, minus stump-to-truck costs and all other relevant on-site costs, such as sale preparation and administration, site preparation and brush disposal, and regeneration. (Neither hauling cost or road construction cost should be deducted from the delivered log value in G_{lit} because these costs are attached to other decision variables.)

$H_{ab,t}^{k}$ = discounted cost of constructing link ab to standard k in period t

$I_{ab,t}^{k,y}$ = discounted cost of a unit of traffic type y traveling from node a to b on road standard k in period t

$I_{ba,t}^{k,y}$ = discounted cost of a unit of traffic type y traveling from node b to a on road standard k in period t.

Other objective functions that may be of interest are to minimize total discounted cost (while meeting timber harvest targets), to maximize timber harvest over the planning horizon (while meeting environmental objectives), or to minimize environmental impacts (while meeting timber harvest targets).

THE HEURISTIC INTEGER PROGRAMMING (HIP) PROCEDURE

The heuristic integer procedure (HIP) described below attempts to find the mix of resource projects, road construction and reconstruction projects, and traffic routing that results in the best objective function value, while satisfying the constraints that have been imposed. Both maximization and minimization problems can be solved. All road construction and reconstruction projects are set to values of 0 or 1 by the procedure. The resource projects are handled as continuous decision variables with an upper bound of 1.

The HIP procedure is an iterative process that utilizes a continuous LP solution in each iteration. First, the mathematical model is reformulated as a continuous LP problem, and additionally, "deviation variables" are added to the side constraints to avoid infeasibilities that may result from HIP rounding decisions. Next, this reformulated problem is solved as a continuous LP problem. The solution is processed by the HIP procedure, essentially a set of decision rules based on the IRPM-type formulation, to determine which road projects to round to values of 0 or 1. These decisions are incorporated into the LP matrix, and another continuous LP solution is obtained. The new solution is processed, and the resulting decisions are incorporated into the LP matrix, and so on. The process continues until all the road projects have been assigned a value of 0 or 1, and no further modifications can be found that improve the objective function value without violating one or more constraints. This generally requires from seven to 12 iterations. The matrix modifications and decision rules are described below.

Formulation Modifications

Several modifications are made to the mathematical formulation described earlier to prepare a linear programming matrix for the HIP solution procedure. First, the road construction and reconstruction projects are defined as continuous decision variables that can assume any nonnegative value. This is necessary to permit solving a model as a linear programming problem.

Second, a set of "deviation variables" are added to each side constraint. These deviation variables prevent the continuous LP from becoming "infeasible" after variables have been rounded to values of 0 and 1 by the HIP program. (When an LP problem is determined to be infeasible, most LP solvers report only the information for the LP iteration in which the infeasibility was detected. Because this is normally only one of many possible infeasible solutions, these solutions do not contain reliable information for making heuristic decisions.) The formulation for an upper bound (less than or equal to) constraint is:

$$\begin{array}{llll}\text{[Original Maximization Objective Function]} & -\,C_{iL}\,(D_{iL}) & -\,C_{iH}\,(D_{iH}) & \\ \text{[Original Left-Hand Side of Constraint } i\text{]} & -(D_{iL}) & -(D_{iH}) & \leq RHS_i \\ & (D_{iL}) & & \leq B_{iL}\end{array}$$

where

D_{iL} = lower-cost deviation variable for constraint i
D_{iH} = higher-cost deviation variable for constraint i
RHS_i = right-hand-side value for constraint i

C_{iL} = penalty assigned to D_{iL}
C_{iH} = penalty assigned to D_{iH}
B_{iL} = upper bound placed on D_{iL}.

To illustrate, if the original left-hand side of a constraint were to exceed RHS_i, D_{iL} and/or D_{iH} would assume a nonzero value, thus avoiding an infeasibility in the LP solution. The same formulation is used for minimization objective functions, with the exception that the objective function penalties, C_{iL} and C_{iH}, carry positive rather than negative signs.

The penalty for violating a constraint is imposed in a step-wise manner, that is, $C_{iH} > C_{iL}$. The logic is that some amount of constraint violation can usually be tolerated because of the inaccuracies and random nature of some constraint coefficients. The amount that can be tolerated, B_{iL}, is specified by the user (the default is 10 percent of RHS_i). Any constraint violation above this amount is forced into the higher-cost deviation variable, D_{iH}, whose coefficient, C_{iH}, is 1.2 times C_{iL}. When this occurs, the heuristic considers the problem to be infeasible. The heuristic solution output reports the objective function value both with and without the deviation variables.

The lower-cost penalty, C_{iL}, is calculated as follows:

C_{iL} = 1.5 times the larger of: R_{max} or 1

where

R_{max} = the maximum $| c_j/a_{ij} |$ (for all j where $a_{ij} \neq 0$) .

Here, c_j is the objective function coefficient for the jth decision variable, and a_{ij} is the constraint coefficient for the jth variable in constraint i. C_{iL} exceeds the contribution per unit of constraint i that each decision variable makes in the objective function, so that deviation variables are in solution at a nonzero value only when necessary to maintain feasibility in a continuous LP solution. In the absence of interdependence among decision variables this would be accomplished by any C_{iL} value exceeding R_{max}. Relationships such as companion relationships, however, may cause c_j/a_{ij} to underestimate the objective function impact per unit of side constraint i. R_{max} is multiplied by 1.5 to allow for such relationships.

If there are no decision variables for which both c_j and a_{ij} are nonzero, then C_{iL} is calculated as:

C_{iL} = 1.5 times the larger of: S_{max} or 1

where

S_{max} = the maximum $| c_j |$ (over all j).

Low- and high-cost deviation variables are also added to lower bound (greater than or equal to) constraints. The only difference is that positive signs are placed on the constraint coefficients for the deviation variables. Four deviation variables are added for equality constraints, a pair of lower and higher cost variables for exceeding the RHS value and a pair for underachieving the RHS value. Four deviation variables are also added when a side constraint contains both an upper and lower bound (called ranges).

After the formulation modifications have been made, the first continuous linear programming solution is obtained.

Iteration 1: Actions Following the First LP Solution

Iteration 1 is a capacity-modification iteration. The coefficients associated with the road projects in the capacity equations (equation 3 in the general formulation) are lowered to create larger fractional values for the road projects

this ratio is small. As a result, something less than the full cost of constructing link *ab* is accounted for in the continuous LP solution. This is often the case in the first LP solution because initially (as discussed earlier) these coefficients represent the maximum traffic capacity.

The capacity coefficients are reduced to more closely approximate the observed traffic volume on the network. As a result, road construction projects enter into subsequent LP solutions at larger fractional values, so that more of the road construction costs are included. This provides improved LP solutions on which to base future heuristic decisions. Once a decision has been made to round a road construction project for link *ab* to 1, the original capacity coefficients for that link are restored to limit traffic to the intended capacity. (See appendix A for a detailed presentation of the capacity adjustment computations.)

Other LP Matrix Modifications—Two additional changes are made to the LP matrix prior to obtaining the next LP solution. First, the constraints that set the sum of the road projects for a link to less than or equal to 1 (equation 1) are made into free (nonconstraining) rows. This permits road construction projects to assume values greater than 1, a change made necessary by the capacity adjustment discussed above. When the capacity coefficients are reduced to increase the fractional value of road projects in solution, there is a chance that the volume of traffic to be routed over a link in a later solution will exceed the new (smaller) capacity. Freeing these constraints allows this to occur.

Second, the objective function coefficients for the traffic variables are divided by 100. The rationale for this modification is: In the continuous LP solutions, two types of transportation costs are considered in routing traffic: (1) the cost associated with the flow variables (haul cost), and (2) the cost associated with road construction or reconstruction. The continuous LP formulation does a very good job of routing traffic so as to minimize haul cost, because the traffic variables were originally formulated as continuous decision variables. But for reasons discussed under capacity adjustment, road construction costs may not be included in their entirety (because the road project is in the solution at a fractional value). As a result, in the continuous LP solutions more weight is placed on haul costs than road construction costs in routing traffic. This is offset in the following LP solution by reducing the objective function coefficients for the flow variables. After the second LP solution is obtained, the original objective function coefficients are restored. But the information about the least-road construction cost routes is captured via another capacity modification in the next iteration.

Iteration 2: Actions Following the Second LP Solution

The primary action taken in this iteration is a second capacity adjustment. The purpose of this adjustment is to further modify capacity on the links to better approximate the actual traffic volumes, while at the same time encouraging future traffic over links carrying relatively large traffic volumes in the current LP solution.

Adjustment of Capacity Coefficients—The capacity for link *ab* is modified to the larger of: (1) the average observed traffic flow on the network, or (2) the observed traffic flow on link *ab*. This adjustment encourages traffic in subsequent LP solutions to travel over the links carrying relatively high volumes of traffic in the current LP solution where the effects of road construction

in the subsequent continuous LP solutions. These larger values for road projects provide an improved basis from which to make rounding decisions. After a road project has been rounded to either 0 or 1, the original coefficients are reinstated. In addition, the first LP solution is analyzed in this iteration to see if processing should stop because the mixed-integer problem is either infeasible, unbounded, or that the optimal solution is to do nothing.

Checks for Infeasible, Unbounded, or Null Solutions—If the first continuous LP solution is infeasible, processing is stopped because the mixed-integer problem, which is more restrictive, must also be infeasible. There are two checks for infeasibility. First, the LP solution status is checked to determine if it is "infeasible." Second, the LP solution is checked to determine if a higher-cost deviation variable is in solution with a nonzero value. This occurs when the LP solver does not find a continuous solution within the bounds placed on the side constraints.

If the first continuous LP solution is unbounded, processing is stopped because this indicates an error in the LP matrix. All the decision variables have upper limits imposed by the relationships in the IRPM-type formulation. The road construction projects are limited to the equivalent of one per road link (equation 1); the resource projects are limited to the equivalent of one per polygon (equation 4); and the traffic variables cannot exceed the capacity coefficients in equation 3. Thus, unbounded LP solutions can be obtained only if there is something wrong with the LP matrix.

If no projects were selected in the first continuous LP solution, processing is stopped. In this case the optimal mixed-integer solution is to do nothing.

Adjustment of Capacity Coefficients—If the linear programming solution is feasible, optimal, and nonnull, then the capacity coefficients in the LP matrix are adjusted downward to match the largest observed flow for each road standard.

The capacity coefficients, $C^k_{ab,g}$, are associated with the road construction projects in the capacity equations (equation 3). To understand the logic behind adjusting them, one must first understand the role they play. Initially, these coefficients represent the maximum traffic volume permitted on link *ab* if constructed to road standard *k*. Given the volume of traffic normally found on forest roads, this limit is typically much larger than the actual volume of traffic.

These coefficients play another role in the continuous LP. They determine the amount of road construction, and therefore the road construction cost that is included in a continuous LP solution. This is most easily explained via a simplified example. The capacity equation for link *ab*, having only one timing option and one road standard option, would be written as:

$$1.0\,(Tab1) - CAP\,(ab1) \leq 0$$

where *Tab1* is the traffic volume over link *ab*, *ab1* is the road construction project (formulated as a continuous variable), and *CAP* is the capacity coefficient. Solving for *ab1* yields:

$$ab1 \geq (1.0/CAP)Tab1$$

Since *ab1* imposes a cost in both cost minimization and profit maximization objective functions, the optimization process in the continuous LP will assign a value to *ab1* as small as possible. This value will be *Tab1/CAP*. When capacity is substantially larger than the actual quantity of traffic,

costs are magnified relative to haul costs. (See appendix B for a detailed presentation of the capacity adjustment computations.)

Other LP Matrix Modifications—One additional change is made to the LP matrix prior to obtaining the next LP solution. The original objective function coefficients for the traffic decision variables are reinstated in the LP matrix.

Iteration 3: Actions Following the Third LP Solution

Iteration 3 is the first of the iterations in which rounding decisions are made.

Rounding Procedure: Round Selected Road Projects to 1—The first step in the rounding procedure is to develop a list of candidate road projects for rounding. Each link having flow in the current continuous LP solution provides one candidate project. The project selected has (1) the lowest road standard that will handle the volume of flow over all periods on the link and (2) a period of implementation that corresponds to the first period in which there is a significant amount of traffic on that link (flow on the link exceeds 5 percent of the average flow over all links in that period). In addition, the road project selected must be compatible with any resource projects that may have been fixed to 1 by the user. For example, if a fixed resource project would require (through a resource project-to-road trigger, equation 8) that link *ab* be constructed in period 1, but no flow occurs until period 2, the period 1 road project would become the candidate.

Next, an index for ranking candidate road projects is calculated for each candidate project. Projects are ranked in ascending order by index value. (Projects with smaller index values are preferable to projects with higher index values.) The index is comprised of the sum of the following components:

1. The contribution of the road project to the side constraints relative to the slack (unused space) in the side constraints weighted by 1 minus the activity level (solution value) for the road project in the current LP solution. The greater this component, the larger (less desirable) the index. The weights cause the contribution to the index to be less, the larger the activity level. Thus, the important road projects in the continuous LP (as measured by the magnitude of the activity level) have less penalty for side constraint impact than do road projects in solution at lower fractional values.
2. The amount of flow on link *ab* relative to the average flow on the network in the current LP solution, multiplied by –1. The more negative this component, the more desirable the index. This component uses quantity of traffic to impute a measure of importance for the links.
3. The relative cost of the road project, multiplied by a weight of 1 minus the activity level for the road project in the current LP solution. The greater this component the less desirable the index. As in 1 above, the weight in this component reduces the amount of penalty imposed in the index for projects having larger activity levels. Thus, the more important projects (as measured by the size of their activity level) have less penalty for construction cost.

Next, determine which candidate projects connect to the "developed portion of the road network." The "developed portion of the road network" consists of road projects representing existing roads, plus projects representing road construction or reconstruction that were previously rounded to 1. Only road projects connecting to the "developed portion of the road network" may be rounded to 1. This gives rise to a continuous developed network that connects to the demand nodes.

Last, from the set of road projects connecting to the "developed portion of the road network," round to 1 the project with the best (lowest) index, provided that project is feasible with each of the side constraints. If the project is not feasible, select the connecting project with the next best (lowest) index, providing it is feasible. At most, no more than 50 percent plus one of the original candidate projects can be rounded to 1 in iteration 3 or 4. This is increased to 75 percent plus 1 beginning with iteration 5. These cut-off points are a compromise between making as many decisions as possible in an iteration to reduce the total number of iterations and avoiding making too many decisions which, as experience has shown, can result in poor mixed-integer solutions. (See appendix C for a detailed presentation of the rounding procedure.)

The feasibility of rounding a candidate road project with regard to side constraint i is checked by adding the impact of that road project to the cumulative impact from projects previously rounded in the current iteration, and comparing that total to the "allotted space" in side constraint i. The "allotted space" is total amount of impact from rounding projects that is permitted in the current iteration. If the "allotted space" is not exceeded, the project is feasible.

Calculation of "allotted space" begins with the "actual space," which is an estimate of the actual space available in a side constraint for absorbing impacts from rounding fractional road projects to 1. The "actual space" is calculated immediately following an LP solution. First, the total impact from the previously rounded road projects, the traffic variables, and the resource projects is calculated by subtracting the impacts of the fractional road projects from the activity level for side constraint i in the LP solution. Next, this total impact is subtracted from the right-hand-side value of the constraint to compute the estimated "actual space."

The "allotted space" is then computed by adjusting "actual space" by an adjustment percentage. This percentage is composed of two factors multiplied together. The first factor varies the direction and amount of adjustment based on the relative importance of road construction in side constraint i. When road construction comprises the majority of the impacts in side constraint i, this factor would reduce "actual space" by approximately 50 percent. The amount of reduction becomes less as the relative importance of road construction in side constraint decreases, reaching zero when roads comprise 10 percent of the side constraint impact. These reductions are designed to avoid committing too much of the available space in a side constraint in any one iteration when roads play an important role in the side constraint. (Experience has shown that objective function value of the resulting HIP solution suffers when a large portion of the rounding decisions are made in any one iteration.) When roads comprise less than 10 percent of the side constraint impact, this factor applies a percentage increase to "actual space." The maximum upward adjustment of 10 percent is reached when roads comprise less than approximately 8 percent of the side-constraint impact. In these cases, roads play such a minor role that impacts exceeding "actual space" can be accommodated in subsequent LP solutions by adjustments in the resource projects and traffic flow variables selected.

The second factor phases out the first factor as the number of iterations increase so that the "actual space" in the side constraints eventually becomes available. The reduction of the first factor begins in iteration 4, the second rounding iteration, and increases in subsequent iterations until it is

completely phased out in iteration 8. Thus, beginning with iteration 8, "allotted space" equals "actual space." (See appendix D for a detailed presentation of the feasibility computations.)

Other LP Matrix Modifications—Two additional changes are made to the LP matrix prior to obtaining the next LP solution. First, the original capacity coefficients are reinstated in the matrix for those links having a road construction project rounded to 1 this iteration. Second, the mutually exclusive constraints (equation 1) associated with the links having a project rounded to 1 are reinstated as less than constraints (they were changed to nonconstraining rows in the first iteration).

Iteration 4: Actions Following the Fourth LP Solution

In iteration 4 and all subsequent iterations, the traffic volumes on the links that contain fractional road projects are compared with the current capacity coefficients on those links. If the difference exceeds a threshold value, this iteration becomes exclusively a capacity adjustment iteration. If the threshold is not exceeded, this iteration reviews the rounding decisions made in previous iterations, and then proceeds to the procedure for rounding fractional road projects.

Adjust Capacity Coefficients if Needed—Beginning with iteration 4, the observed flow on the links having fractional road projects is compared with the current capacity on those links. If the observed flows are, on the average, less than 10 percent of capacity for the standard that matches the observed quantity of traffic (the lowest road standard that can accommodate the flow on those links over all periods) and capacity was not adjusted in the previous iteration, a capacity adjustment is made. This adjustment is identical to the iteration 2 procedure and is described in detail in appendix B.

It may not be obvious why further capacity adjustments are necessary. The capacity adjustment in iteration 2 is based on the average traffic flow over links having fractional road projects at that time. As the rounding iterations proceed, the road projects rounded first are those closest to the demand nodes carrying the larger quantities of traffic. The remaining links having fractional road projects tend to be in the smaller tributaries of the network and carry a smaller amount of traffic. If the flow on these links is small relative to capacity, it is advantageous to adjust capacity downward for the reasons described earlier—larger fractional values for road projects in later iterations, which provide a better basis for rounding projects.

If capacity for the fractional links is adjusted, the modified capacity coefficients are incorporated into the LP matrix, and another LP solution is obtained for the next iteration. Alternatively, if the observed flows do exceed 10 percent of capacity on the average, no capacity adjustment is undertaken and the iteration proceeds as described below.

Change Road Standard Selections?—Road projects rounded to 1 in previous iterations are reviewed to determine if they should be replaced with projects having either a lower or higher road standard. The criteria are as follows:

1. Lower the road standard for link *ab* if the largest traffic flow in any one period is less than the capacity of a lower road standard in that period. Tentatively select the lowest standard whose capacity exceeds the observed traffic flow on link *ab* in each period. If the lower-standard project is feasible with regard to the side constraints, round it to 1 and fix to 0 all other

projects defined for that link. If the tentatively selected project is infeasible, do not make any change in standard.

2. Raise the road standard selection for link *ab* if the traffic flow in any period equals the capacity for the road project previously rounded to 1. Tentatively select the next higher standard. If that project is feasible with regard to the side constraints, round it to 1 and fix to 0 all other projects defined for that link. If the tentatively selected project is infeasible, do not make any change in standard.

Rounding Procedure: Round Selected Road Projects to 1—If modifications were made to "the developed portion of the road network" during this iteration (for example, change in road standard), and those modifications resulted in a critical change, then this iteration stops without rounding any additional fractional projects. A change is considered critical if:

1. Any side constraint that was within 10 percent of its bounds before modifying "the developed portion of the road network" is no longer within 10 percent of its bounds after the modifications.
2. Any side constraint that was not within 10 percent of its bounds before modifying "the developed portion of the road network" is now within 10 percent of its bounds after the modifications.

(See appendix D for a detailed presentation of this feasibility testing.) Modifications in "the developed portion of the network" resulting in critical changes can have a significant effect on the selection of resource projects, fractional road projects, and traffic routing. So, if either of these conditions are present, the LP is solved again and rounding decisions are postponed.

If modifications in the road standard selections do not produce a critical change, the rounding procedure described earlier for iteration 3 is undertaken.

Other LP Matrix Modifications—These are the same as the changes made in iteration 3.

Iteration 5: Actions Following the Fifth LP Solution

Actions taken in iteration 5 are identical to iteration 4 with the exception of an additional check to determine if the timing for construction should be postponed to a later period.

Adjust Capacity Coefficients if Needed—This step is the same as in iteration 4.

Change Road Standard Selections?—This step is the same as in iteration 4.

Postpone Road Construction?—Road projects rounded to 1 in previous iterations are reviewed to determine if they should be replaced with projects implementing road construction in a later period. Postponing construction is attempted if in the period of construction the amount of traffic on link *ab* is less than 5 percent of the average flow on the links in the network. Tentatively select the project of the same standard whose construction period coincides with the earliest period in which flow exceeds 5 percent of the average network flow. If this project is feasible with regard to the side constraints, round it to 1, and set all other projects for the link to 0. If this project is infeasible with regard to the side constraints, make no change. If there is no later period in which flow exceeds 5 percent of the average

network flow, no change is made in the postponing routine. This situation, however, is addressed in the "close links" routine beginning in iteration 6.

Rounding Procedure: Round Selected Road Projects to 1—This step is the same as in iteration 4.

Other LP Matrix Modifications—These are the same as the changes made in iteration 3.

Iterations 6+: Actions Following the Sixth to Last LP Solution

This iteration is identical to iteration 5 with the exception of an additional check to determine if any link previously selected for road construction should be closed.

Adjust Capacity Coefficients if Needed—This step is the same as in iteration 4.

Change Road Standard Selections?—This step is the same as in iteration 4.

Postpone Road Construction?—This step is the same as in iteration 5.

Close Links?—Road projects rounded to 1 in previous iterations are reviewed to determine if sufficient traffic is being carried to warrant the road project. If the amount of traffic on link *ab* in each period is less than 5 percent of the average flow on the links in the network, that link is closed. This is accomplished by setting all road projects pertaining to that link to zero in the LP matrix.

Rounding Procedure: Round Selected Road Projects to 1—This step is done the same as in iteration 4.

Other LP Matrix Modifications—These changes are the same as those made in iteration 3.

Stopping Rule—The solution process stops when the current LP solution contains integer values for all road projects and there are no desirable changes to be made to previous rounding decisions (changing standard, postponing, or closing links). This generally occurs within seven to 12 iterations. The final LP solution is the integer solution for roads having the best objective function value, while satisfying the side constraints. In nearly all cases this is the last LP solution made in the heuristic process.

It is possible that the heuristic solution process may not be able to find a feasible mixed-integer solution (all higher-cost deviation variables equal 0) even though a continuous optimal solution was found in the first iteration. This could be because there is no feasible mixed-integer solution (mixed-integer solutions are more restrictive than continuous solutions). Or it may be that a mixed-integer optimal exists but was simply not found by the heuristic procedure. In any case, when this occurs the solution process reports the best mixed-integer solution it can find. The value assigned to the deviation variables measures the extent to which side constraints are violated.

SOLUTION VALIDATION

Comparison With Optimal Mixed-Integer Solutions

One method to validate HIP solutions is to compare them with the optimal mixed-integer solutions obtained using the branch-and-bound approach. Doing this on a large scale, however, proved to be cost-prohibitive. This comparison was made with Small Twin Rocks, a land management

and transportation planning problem adapted from a portion of the Twin Rocks area (Jones and others 1986). The size statistics for this model are summarized in table 1. The polygons and road network are presented in figures 3 and 4, respectively.

Table 1—Size statistics for the five models included in this report

Category	Small Twin Rocks	Copeland Creek	Moose Creek	Lowman area	TUJO-RECOYLE area
Polygons	28	282	301	324	569
Road Links:					
Proposed	22	84	194	126	167
Existing	2	64	79	104	67
Total	24	148	273	231	234
Demand nodes	2	2	4	4	1
Periods	3	3	2	3	3
Matrix data:					
Rows	365	2,281	2,150	3,326	3,033
Columns	260	1,651	1,923	3,515	4,046
Integers	67	318	546	693	702
Density[1]	3.1%	0.8%	0.4%	0.4%	0.4%

[1]Percentage of matrix coefficients that are nonzero.

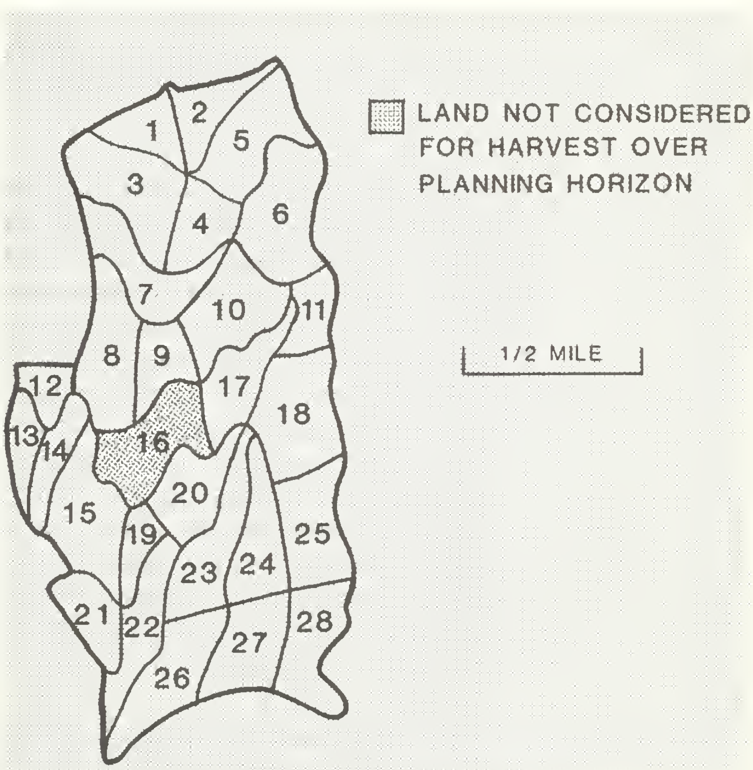

Figure 3—Polygons for the Small Twin Rocks planning model.

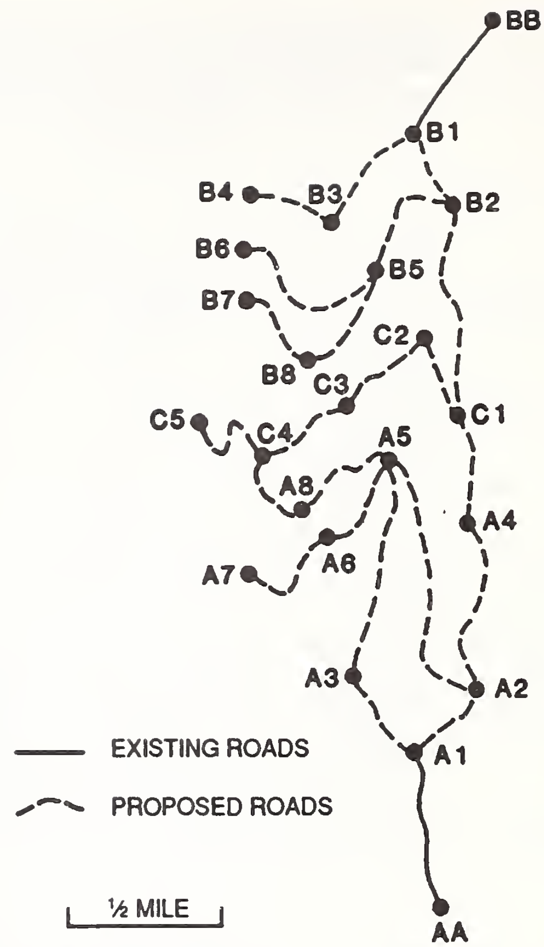

Figure 4—Road network for the Small Twin Rocks planning model.

Six different formulations of this model were solved with both the HIP process and the branch-and-bound algorithm in the IBM Mathematical Programming System (MPSX). These formulations represent a combination of constraints and objectives typical of an actual planning problem.

Table 2 summarizes the six formulations and the resulting objective function values for: (1) the first continuous LP solution in the HIP process, (2) the HIP solution, and (3) the optimal mixed-integer solution. The objective in formulations 1 through 4 was to maximize present net value (PNV). The side constraints in these formulations varied from none (formulation 1), to upper bounds of 450 tons of additional sediment production in each period and lower bounds of 3,000 M bd ft of timber harvest in each period (formulation 4). The objective in formulation 5 was to minimize the production of additional sediment while harvesting at least 3,000 M bd ft of timber in each period. The objective in formulation 6 was to minimize discounted total cost while meeting the same timber harvest targets as formulation 5.

The objective values of the HIP solutions were within 10 percent of the value of the optimal mixed-integer solution in all but one case, formulation 3. Thus, although the HIP process did not find optimal mixed-integer solutions, it did find solutions whose objective function values closely approximated the values for the optimal mixed-integer solutions.

Table 2—Comparing HIP solutions with the optimal mixed-integer solutions for six formulations of the Small Twin Rocks test model

Category	Formulation					
	1	2	3	4	5	6
Objective	Maximize PNV[1]	Maximize PNV	Maximize PNV	Maximize PNV	Minimize SED1[2]	Minimize DTCOST[3]
Side constraints	None	SED1 ≤ 300	SED1 ≤ 300 SED2[5] ≤ 300 SED3[7] ≤ 300	SED1 ≤ 450 SED2 ≤ 450 SED3 ≤ 450 TIMBER1 ≥ 3,000 TIMBER2 ≥ 3,000 TIMBER3 ≥ 3,000	TIMBER1[4] ≥ 3,000 TIMBER2[6] ≥ 3,000 TIMBER3[8] ≥ 3,000	TIMBER1 ≥ 3,000 TIMBER2 ≥ 3,000 TIMBER3 ≥ 3,000
Objective function values:						
(1) First continuous LP	175.2	175.2	175.2	175.2	131.3	231.0
(2) HIP	80.8	69.1	61.2	77.6	313.0	335.7
(3) Optimal mixed-integer solution	85.1	71.5	71.5	79.6	289.5	319.7
Difference between (3) and (2) as a percentage of (3)	5.0	3.4	14.4	2.5	8.1	5.0

[1]PNV = Present net value (thousands of dollars).
[2]SED1 = Additional sediment production in period 1 (tons).
[3]DTCOST = Discounted total cost (thousands of dollars).
[4]TIMBER1 = Timber harvest in period 1 (thousands of board feet).
[5]SED2 = Additional sediment production in period 2 (tons).
[6]TIMBER2 = Timber harvest in period 2 (thousands of board feet).
[7]SED3 = Additional sediment production in period 3 (tons).
[8]TIMBER3 = Timber harvest in period 3 (thousands of board feet).

Comparison With Other Planning Methods

Plans developed via current methods for tactical planning provide, perhaps, a more useful basis for comparison. Unlike comparisons with optimal mixed-integer solutions, these comparisons can be done with full-scale areas. They also measure the potential gains (or losses) in efficiency that can be expected with the HIP solution procedure.

A common method for analyzing when and where to harvest timber and construct roads in the Northern Rockies utilizes professional judgment to make the harvesting selections and a road network cost minimization model, such as NETWORK (Sessions n.d.), to make the roading selections. Plans developed via the current HIP algorithm were compared against this method on two planning areas in the Northern Rockies. The Copeland Creek area, on the Kootenai National Forest in Montana, was one of the areas included in an earlier study (Jones and others 1986). For this area, the HIP process developed a new management alternative, which was compared to a management alternative developed in that earlier study using the conventional approach. The second area, the Moose Creek area, is on the Helena National Forest. The model for this area was developed, solved, and the results compared by National Forest System personnel (Bower 1990). For both areas, identical data and management objectives were used in the conventional and HIP approach.

The results of these comparisons are summarized in table 3. The objective for the Copeland Creek area was to find the combination of harvesting and road building that maximizes present net value, while maintaining the additional sedimentation in each of four watersheds below management-specified upper bounds. The objective for the Moose Creek area was to find the combination of harvesting and road building that maximizes present net value, without exceeding upper bounds on the number of acres having regeneration harvests.

Table 3—Comparing HIP solutions with plans developed via a process in which managers choose the arrangement and timing of land management activities and then use a network cost minimization model to find the least-cost access routes

Category	Copeland Creek area	Moose Creek area
Objective	Maximize PNV[1]	Maximize PNV
Side constraints	Additional sediment in each of four drainages must not exceed specified amounts in each period	The number of acres having regeneration harvests must not exceed the specified amounts in each period
Objective function values:		
(1) First continuous LP	4,491	1,665
(2) HIP	4,296	1,332
(3) Manager choice with road cost minimization	3,014	978
Difference between (3) and (2) as a percentage of (3)	43	35

[1]Present net value (thousands of dollars).

In both cases the HIP solution process developed plans that met the constraints and with objective function values that substantially exceeded the professional judgment/network cost minimization approach. For Copeland Creek, the increase was 43 percent; for Moose Creek, 35 percent. For the Copeland area, the objective value for the HIP plan was within 5 percent of the value for the first continuous LP, indicating the objective value of this solution closely approximates the optimal mixed-integer value. For Moose Creek the objective value was about 21 percent less than the value for the first continuous LP solution. These results are of approximately the same magnitude as comparisons made in an earlier study (Jones and others 1986).

Comparison With the First Continuous LP Solution

The first continuous LP solution made in the HIP process provides another basis for comparison. This LP solution is the optimal solution to the continuous version of a model (road projects can assume any fractional value). Restricting road projects to integer values places an additional limitation on a model. This additional limitation will decrease the objective function value for maximization problems, increase it for minimization problems, or possibly have no effect (all road projects have integer values in the first continuous LP solution). An additional limitation can never improve the objective function value of a continuous LP problem. The objective function value of the first continuous LP is, therefore, an upper bound for the objective function value of the optimal mixed-integer solution.

This upper bound provides a convenient, but imperfect yardstick. It is imperfect because it is unknown how tight that bound is for any one problem (unless, of course, the optimal mixed-integer solution is known). Useful information is obtained by this comparison only when the difference between the HIP solution and the first LP solution is small. A large difference does not necessarily indicate a poor HIP solution, because it could be reflecting a large difference between the optimal MIP solution and the first LP solution. For example, in the Small Twin Rocks model (table 2) the difference between the HIP and first LP solution was quite large, while the difference between the HIP and optimal MIP solution was small.

Table 4 compares the objective function values of 15 HIP solutions with the respective values from the first continuous LP solutions. These solutions represent variations in the constraints, level of constraints, and objective functions for three full-scale models. One was the Moose Creek model described earlier. The second model was developed to analyze salvage harvesting options for the Lowman Fire Complex on the Boise National Forest in Idaho (Bower 1990). The third model was developed for an area on the Lolo National Forest in Montana called TUJO-RECOYLE (Bower 1990). The size statistics for these models are presented in table 1.

The HIP objective function values for the Lowman model were within 5 percent of the value of the first LP solution in all but one case, and within 10 percent in all cases. For these cases, the objective function values of the HIP solution closely approximated the mixed-integer optimums. For the Moose Creek and TUJO-RECOYLE models, the differences between the objective function value of the first continuous LP solution and the HIP solution were generally somewhat larger. For Moose Creek the deviations varied from 24 percent to 29 percent, while for the TUJO-RECOYLE model they varied from 4 percent to 21 percent. These larger deviations, however, do not necessarily indicate poor solutions for the reasons discussed earlier.

Table 4—Comparing the objective values of HIP solutions made for full-scale models with the objective values of the first continuous LP solutions made in the HIP process. (Models vary by the objective function, and by the level and type of side constraints)

Solution number	Model	Objective	Objective values: First LP	Objective values: HIP	Relative deviation[1]
			- - *Thousand dollars* - -		*Percent*
1	Lowman area	Max. PNV[2]	1,686	1,529	9.3
2	Lowman area	Max. PNV	6,381	6,284	1.5
3	Lowman area	Max. PNV	5,788	5,668	2.2
4	Lowman area	Max. PNV	4,602	4,512	1.9
5	Lowman area	Max. PNV	5,308	5,188	2.3
6	Lowman area	Max. PNV	4,896	4,777	2.4
7	Moose Creek	Max. PNV	2,056	1,504	26.8
8	Moose Creek	Max. PNV	1,665	1,227	26.3
9	Moose Creek	Max. PNV	1,579	1,198	24.1
10	Moose Creek	Max. PNV	1,932	1,366	29.3
11	TUJO-RECOYLE	Max. PNV	–3,802	–4,219	11.0
12	TUJO-RECOYLE	Max. VQI[3]	2,616	2,765	5.7
13	TUJO-RECOYLE	Max. PNV	–2,038	–2,276	11.7
14	TUJO-RECOYLE	Max. VOL[4]	110	105	4.1
15	TUJO-RECOYLE	Max. VOL	110	86	21.5

[1]The difference in objective function values divided by the value for the first LP, multiplied by 100 to express the relative deviation as a percent.
[2]PNV is present net value.
[3]VQI is a visual quality index applying over all periods.
[4]VOL is timber harvest volume over all periods.

Computer Requirements

The 15 full-scale model solutions presented in table 4 provide a good basis for estimating the computer requirements of the HIP solution process. The computation times and costs associated with these solutions are summarized in table 5. Perhaps the most relevant categories are the CPU and I/O times because they have the most bearing in predicting performance on other computer systems. (Resource time is actually a sum of the other time components multiplied by the amount of memory used, and CC/ER is also based on CPU time.) CPU time averaged 21.9 minutes and I/O time averaged 47.9 minutes. The average total cost of these solutions, based on the

Table 5—Computer times and costs for the 15 full-scale model HIP solutions reported in table 3 run on System C at NCC-FC (UNISYS Model 1193)

	NCC-FC cost categories: Resource time[1]	CPU time[2]	I/O time[3]	CC/ER time[4]	Total cost[5]
	- - - - - - *Minutes*	- - - - - -	- - - - - -	- - - - - -	
Average	242.0	21.9	47.9	10.7	$120.07
Highest	349.2	37.0	74.1	12.6	198.67
Lowest	163.6	13.5	22.5	8.4	57.37

[1]Resource time, actually an index for assessing a memory charge, is the sum of CPU, I/O, and CC/ER time multiplied by memory usage.
[2]CPU time is central processing unit time.
[3]I/O time is time devoted to input-output devices.
[4]CC/ER is a table look-up value based on approximate CPU time for a typical user request.
[5]Total cost represents the composite measured in dollars using the July 1990 price structure for t-priority at NCC-FC.

July 1990 T-priority rates on System C at NCC-FC, was $120 and ranged from $57 to $199. While not small, these costs are modest when compared to the cost of solving these models with the mixed-integer branch-and-bound algorithm available in the Functional Mathematical Programming System (Sperry Corporation 1984) available at NCC-FC. Earlier attempts at solving models of a similar size using this branch-and-bound algorithm typically cost in excess of $500 and still did not find the optimal, or in some cases even a feasible, mixed-integer solution.

Below, the cost of a HIP solution is expressed as a multiple of the cost of the first continuous LP solution made in the HIP process. This solution is always made from the slack (beginning) basis. Estimates of the resources required to run the HIP procedure on different computers can be made by multiplying these factors by the time required to make a continuous LP solution of an IRPM-type formulation on that computer. The multipliers are:

Category	**HIP procedure as a multiple of the first LP solution**
Resource time	13.7
CPU time	8.5
I/O time	16.2
CC/ER time	24.2
Total cost	11.6

In the current configuration, the LP solver and the program containing the heuristic decision rules are written as separate programs. They are linked together in a control language loop. This means that each time these processes are accessed (typically seven to 12 times for a HIP solution) the internal arrays and variables must be reinitialized from data stored in files. In addition, the two processes communicate with each other via files. There is a potential for substantial savings in I/O-related costs by integrating the decision rule program and an LP solver into one program.

CONCLUSIONS

The HIP solution procedure was designed to provide efficient, feasible mixed-integer solutions for the IRPM-type formulation at a reasonable cost. To this point we have had good success in achieving feasible mixed-integer solutions for both test models and full-scale planning models. Costs for full-scale models can be expected to fall in the range of $50-$200 in T-priority, given July 1990 rates at NCC-FC. These costs are modest compared to the costs associated with achieving mixed-integer solutions to these models using branch-and-bound algorithms.

Comparisons made with optimal mixed-integer solutions and other planning methods have been encouraging. Plans developed using the HIP solution procedure have been found to increase present net value objectives by as much as 40 percent over plans developed via a planning approach frequently used in the Northern Rockies. In comparisons made with optimal mixed-integer solutions, the objective function value for the HIP solutions was within 10 percent of the optimal value in five of six cases.

We believe the HIP procedure has much potential for arranging and scheduling land management activities and road construction projects in ways that efficiently meet specified management objectives. This efficiency should be valuable in helping land managers meet the many demands that are placed on our forest lands.

REFERENCES

Bower, Fred. 1990. [Personal communication]. May 14. Missoula, MT: U.S. Department of Agriculture, Forest Service, Northern Region.

Johnson, K. Norman; Jones, Daniel B. 1979. A user's guide to multiple use sustained yield resource scheduling calculation (MUSYC). Fort Collins, CO: U.S. Department of Agriculture, Forest Service, Timber Management. 242 p.

Johnson, K. Norman; Stuart, Thomas W.; Crim, Sarah A. 1986. FORPLAN Version 2: an overview. Washington, DC: U.S. Department of Agriculture, Forest Service, Land Management Planning Systems Section. n.p.

Jones, J. Greg; Hyde, James F. C., III; Meacham, Mary L. 1986. Four analytical approaches for integrating land management and transportation planning on forest lands. Res. Pap. INT-361. Ogden, UT: U.S. Department of Agriculture, Forest Service, Intermountain Research Station. 33 p.

Jones, J. Greg; Meacham, Mary L.; Weintraub, Andres; Magendzo, Adrian; Cox, Wallace R.; Bower, Fred; Bergan, Ernest E.; Kirby, Malcolm W. 1990. IRAS: a system for site-specific land management and transportation planning. In: Proceedings of the 1989 Society of American Foresters National Convention; 1989 September 24-27; Spokane, WA. Washington, DC: Society of American Foresters: 355-360.

Kirby, Malcolm W.; Wong, Peter; Hager, William A.; Huddleston, May E. 1980. Guide to the Integrated Resource Planning Model. Berkeley, CA: U.S. Department of Agriculture, Management Sciences Staff. 212 p.

Navon, Daniel I. 1971. Timber RAM...a long-range planning method for commercial timber lands under multiple-use management. Res. Pap. PSW-70. Berkeley, CA: U.S. Department of Agriculture, Forest Service, Pacific Southwest Forest and Range Experiment Station. 22 p.

Sessions, John. [n.d.] Network analysis using micro-computers for logging planning. Corvallis, OR: Oregon State University, College of Forestry, Department of Forest Engineering. 5 p.

Sperry Corporation. 1984. Functional mathematical programming system, FMPS level 9R2, programmer reference. Sperry Corporation. n.p.

U.S. Department of Agriculture, Forest Service. 1975. Resource allocation analysis. Resource capability system, part VII. Fort Collins, CO: Watershed Systems Development Unit. n.p.

Weintraub, A. 1982. A heuristic approach for solving forest management transportation problems. Serie Documentos de Trabajo No. 82/05/C. Santiago, Chile: Universidad de Chile, Facultad de Ciencias Fisicas y Matematicas, Departamento de Ingenieria Industrial. 19 p.

APPENDIX A: CAPACITY ADJUSTMENT NO. 1

This capacity adjustment is done in iteration 1. When a problem is formulated, it is typical for the capacity of the roads to greatly exceed the observed quantity of traffic on those roads. As explained in the main text under iteration 1, this results in small fractional values being assigned to the road projects in the first continuous LP solution. The objective of this capacity adjustment is to perform a general lowering of the capacity coefficients so that they more closely match the quantities of traffic observed on the road network in the first LP solution.

Capacity adjustment No. 1 is made as follows:

$$MC^{k}_{ab,t} = C^{k-1}_{ab,t} + AF1^{k}\,(C^{k}_{ab,t} - C^{k-1}_{ab,t})$$

where $MC^{k}_{ab,t}$ is the new capacity coefficient to be incorporated into the LP matrix, $C^{k}_{ab,t}$ is the current capacity coefficient for link ab in period t, $C^{k-1}_{ab,t}$ is the current capacity coefficient for the road standard immediately below k (if k is the lowest road standard, $C^{k-1}_{ab,t}$ equals 0), and $AF1^{k}$ is an adjustment factor which is computed for each road standard.

The first step in computing $AF1^{k}$ is to calculate preliminary adjustment factors, $AF1^{k}_{prelim}$, for each road standard:

$$AF1^{k}_{prelim} = \text{the maximum } (F^{s}_{ab,p}/C^{s}_{ab,p})$$

where

$F^{s}_{ab,p}$ = the largest total amount of traffic (traffic is converted to the units of measure defined for capacity and summed over all traffic types) observed on link ab in any one period. Superscript s identifies the lowest road standard that can handle that volume of traffic, and p is the period in which that traffic was observed.

$C^{s}_{ab,p}$ = the capacity coefficient for the road standard, s, and period, p, associated with $F^{s}_{ab,p}$ on link ab.

Separate adjustment factors, $AF1^{k}$, are then computed for each k road standard. For the highest-capacity road standard:

$$AF1^{k} = AF1^{k}_{prelim}$$

providing there was at least some observed traffic on this road standard. If there was no traffic on this road standard (all traffic on the network can be carried by lower standards), $AF1^{k}$ is fixed to 0.15.

For road standards below the highest-capacity road standard, the adjustment factor, $AF1^{k}$, is computed according to the following rules. These rules are based on whether the largest flow, $F^{s}_{ab,p}$, matches the road standard being processed, standard k. $F^{s}_{ab,p}$ is said to match road standard k if the capacity of standard k exceeds $F^{s}_{ab,p}$, while the capacity for the next lower road standard $(k-1)$ is less than $F^{s}_{ab,p}$.

1. If the largest flow, $F^{s}_{ab,p}$, matches the road standard k $(k = s)$ on at least 5 percent of the network links, then:

   $AF1^{k}$ = the larger of $AF1^{k}_{prelim}$ or 0.3

   (The minimum of 0.3 prevents excessive modification.)

2. If the largest flow, $F^{s}_{ab,p}$, matches road standard k ($k = s$) on less than 5 percent (or none) of the network links, and 95 percent or more of the flow on the network is on lower standards, then:

   $AF1^{k} = 0.15$

3. If the largest flow, $F^{s}_{ab,p}$, matches road standard k ($k = s$) on less than 5 percent (or none) of the network links, but less than 95 percent of the network flow is on lower standards, then:

   $AF1^{k} = 1$

   (In this case no capacity adjustment is made because there is insufficient information for modifying capacity.)

APPENDIX B: CAPACITY ADJUSTMENT NO. 2

This capacity adjustment is made in iteration 2, and in later iterations (beginning with iteration 4) when there is a need for an additional capacity adjustment. The objective of this adjustment is to further reduce the capacity coefficients associated with fractional links to promote larger fractional values in subsequent iterations. But care must be taken to avoid reducing capacity so far that traffic is redirected to inefficient routes in the later iterations. This is accomplished by adjusting capacity for the road standard matching the quantity of flow in the current LP solution to the larger of: (1) the average traffic flow on the network links or (2) the observed flow on the link being processed.

Capacity adjustment No. 2 is based on $AF2$, a factor that applies across all road standards in a model. The first step in computing $AF2$ is to compute $F^{s}_{ab,p}$ for each link:

$F^{s}_{ab,p}$ = the largest total amount of traffic (traffic is converted to the units of measure defined for capacity and summed over all traffic types) observed on link ab in any one period. Superscript s identifies the lowest road standard that can handle that quantity of traffic, and p is the period in which that traffic was observed.

The adjustment factor, $AF2$, is computed as:

$$AF2 = \text{the larger of } \{0.3, \text{ or } [\sum_{ab} (F^{s}_{ab,p}/C^{s}_{ab,p})]/N\}$$

where $C^{s}_{ab,p}$ is the capacity coefficient for the road standard matching the largest flow on link ab, which occurs in period p, and N is the number of links for which a value of $F^{s}_{ab,p}$ was computed (all links having at least some traffic are included).

The rules for applying the capacity adjustment factor $AF2$ to link ab vary depending on whether the road standard to be adjusted equals, is lower, or is higher than the road standard matching the highest traffic volume, $F^{s}_{ab,p}$, on the link. The matching road standard is the lowest road standard that will handle the highest traffic volume on the link.

1. If the road standard being adjusted is standard s, the matching standard for the link, capacity for each period (t) is adjusted according to the following rules:

   (i) If $(AF2)(C^{s}_{ab,p}) \geq F^{s}_{ab,p}$, then the capacity for standard s is adjusted for each period as follows:

   $$MC^{s}_{ab,t} = (AF2)(C^{s}_{ab,t}) \qquad \text{(for each } t\text{)}$$

   (ii) If $(AF2)(C^{s}_{ab,p}) < F^{s}_{ab,p}$, then the adjustment for each period is:

   $$MC^{s}_{ab,t} = (F^{s}_{ab,p})(C^{s}_{ab,t} / C^{s}_{ab,p}) \qquad \text{(for each } t\text{)}$$

   $MC^{s}_{ab,t}$ is the new capacity coefficient to be incorporated into the LP matrix for standard s in period t, $C^{s}_{ab,t}$ is the previous capacity coefficient for standard s in period t (t = 1, 2, etc.), and $C^{s}_{ab,p}$ is the previous capacity coefficient for standard s in period p, the period of highest observed flow on link ab. The ratio $(C^{s}_{ab,t}/C^{s}_{ab,p})$ is included to adjust $F^{s}_{ab,p}$ when the capacity for road standard s varies across planning periods. This occurs when the periods are of unequal length.

2. If the road standard being adjusted is higher than standard s (the lowest standard that can handle the observed flow on link ab) capacity is adjusted as follows:

   $$MC^{k}_{ab,t} = C^{k-1}_{ab,t} + (AF2)(C^{k}_{ab,t} - C^{k-1}_{ab,t})$$

   where $C^{k}_{ab,t}$ is the previous capacity coefficient for standard k on link ab in period t, and $C^{k-1}_{ab,t}$ is the previous capacity coefficient in period t for the next lower road standard, $k-1$.

3. If the road standard being adjusted is lower than standard s (the lowest standard that can handle the observed flow on link ab) capacity is not adjusted. There is no reason to adjust capacity downward because the flow on link ab exceeds the capacity for this road standard.

4. If link ab carries no traffic in the current LP solution, the capacity coefficients are modified as follows for all the standards defined for that link:

   $$MC^{k}_{ab,t} = C^{k-1}_{ab,t} + (AF2)(C^{k}_{ab,t} - C^{k-1}_{ab,t})$$

   where all terms are as previously defined. If k is the lowest road standard defined for link ab, then $C^{k-1}_{ab,t}$ equals 0.

APPENDIX C: ROUNDING PROCEDURE

The rounding procedure consists of three major steps: (1) develop a list of candidate road projects for rounding, (2) compute an index for ranking the candidate projects, (3) choose the projects to be rounded to 1 from the set of candidate projects connecting to either a final demand node or to road projects that were previously rounded to 1.

Develop a List of Candidate Road Projects for Rounding

One candidate project is selected for each fractional link that carried traffic in the current LP solution, or was associated with a resource project that contributed flow to the road network. (Association with resource projects is through resource project-to-road triggers. See equation 8 in the general formulation.) First, a preferred period and road standard are selected. Then, a candidate project is selected after comparing the preferred period and standard selections with the feasible projects available for that link.

Preferred Period—The preferred period for link *ab*, P_{ab}, is the earliest period in which there is a significant quantity of traffic on link *ab*. P_{ab} is the earliest period in which:

$$T_{ab,t} \geq MFLOW_t$$

where $T_{ab,t}$ is the total traffic on link *ab* in period *t* (each traffic type is converted to capacity units and summed) and $MFLOW_t$ equals the smaller nonzero value of: $MFLOW1_t$ or $MFLOW2_t$. $MFLOW1_t$, a minimum flow criterion based on the average network flow, is calculated as:

$$MFLOW1_t = (0.05)\,(\Sigma\, T_{ab,t})/N$$

where $T_{ab,t}$ is as defined above, and *N* is the number of links having nonzero flow in period *t*. $MFLOW2_t$ is a minimum flow criterion based on the flow loaded onto the network by resource projects. For all periods except the last period:

$$MFLOW2_t = (\tfrac{1}{2})(FIXRP_t)$$

where $FIXRP_t$ is the smallest quantity of traffic loaded onto the network in period *t* from any one resource project that was preselected (set into solution) by the user. If a resource project creates flow in more than one period, however, only the initial-period traffic is considered in calculating $MFLOW2_t$. For the last period:

$$MFLOW2_t = (\tfrac{1}{2})(RP_t)$$

where RP_t is the smallest quantity of traffic loaded onto the network in the last period by any resource project (whether fixed by the user or not) whose first traffic loading occurs in the last period.

$MFLOW2_t$ is designed to ensure that period selection is consistent with providing access to necessary resource projects. In all but the last period, necessary projects are those preselected by the user. For the last period, all resource projects are considered necessary to accommodate the transhipment formulation where at least one harvesting option must be selected from each polygon (in the transhipment formulation equation 4 is modified from an upper-bound constraint to an equality).

Preferred Road Standard—The preferred road standard for link *ab*, S_{ab}, is the lowest road standard whose capacity exceeds both (1) the largest total traffic flow observed on that link in any one period and (2) the largest total traffic flow loaded onto the network in any one period from the resource projects associated with link *ab* (through resource project-to-road triggers). These total flows include all traffic types converted to capacity units. If the observed quantity of traffic in any one period exceeds the capacity for the highest capacity road standard defined for that link, the highest capacity standard is selected. (After the road project is rounded to 1, capacity becomes an upper bound on the total quantity of traffic.)

Candidate Selection—The selection of the candidate project for link *ab* is based on the following rules:

1. If the preferred road standard, S_{ab}, represents an existing road requiring no reconstruction (there is no cost associated with the road project), then assign a value of 1 to the road project on link *ab* having that standard in the earliest period available. This road already exists on the ground, and the decision rules indicate that it should not be replaced.
2. If the preferred road standard, S_{ab}, is not an existing road, but is available as a proposed project in the preferred period, P_{ab}, then that project is selected as the candidate for link *ab*, providing that selection is not precluded by a project preselected by a user. Preselected projects include road construction projects for link *ab* and resource projects that trigger other road projects for link *ab* that were incorporated into the solution by the user.
3. If the preferred road standard, S_{ab}, is not an existing road and is not available as a proposed project in the preferred period, P_{ab} (because that project is infeasible as a result of a road or resource project preselected by the user, or because a project of standard S_{ab} was not defined for period P_{ab}), then an alternate project is selected. The first priority is to find a project with a different road standard that has the same implementation period. If the quantity of traffic on that link approaches the capacity of the selected road standard, projects with higher road standards are preferred. Otherwise projects with lower road standards are preferred. If no acceptable project during the same implementation period is found, projects with earlier implementation periods are investigated. There must be at least one project feasible for the observed traffic in the current solution to have occurred. The details of this selection process follow. Let $F_{ab,p}$ represent the largest traffic flow for link *ab*; the subscript *p* indicates the period in which the largest flow occurred. Let $C^{s}_{ab,p}$ represent the capacity for the preferred road standard in period *p*, and $C^{s-1}_{ab,p}$ represent the capacity in period *p* for the next lower road standard.

If: $F_{ab,p} \geq C^{s-1}_{ab,p} + 0.75(C^{s}_{ab,p} - C^{s-1}_{ab,p})$

then select the project with the next higher road standard, providing it is feasible with regard to any road or resource projects associated with link *ab* that were preselected by the user. If no feasible higher standards are available, select the project with next lower standard, providing it is feasible. If no feasible road standards are available in the preferred period, repeat this process for period P_{ab}–1, then P_{ab}–2, and so on. There must be an available road project in one of these periods for traffic to have appeared on link *ab* in time P_{ab}.

If: $F_{ab,p} < C^{s-1}_{ab,p} + 0.75(C^{s}_{ab,p} - C^{s-1}_{ab,p})$

then select the project with the next lower road standard, providing it is feasible with regard to any road or resource projects associated with link *ab* that were preselected by the user. If no feasible lower standards are available, select the project with next higher standard, providing it is feasible. If no feasible road standards are available in the preferred period, repeat this process for period P_{ab}–1, then P_{ab}–2 and so on. There must be an available road project in one of these periods for traffic to have appeared on link *ab* in time P_{ab}.

Calculate Index Values for the Candidate Projects

The next step is to calculate an index for ranking the candidate road projects. The smaller the value of this index, the more preferred the road project. The index for road project *j* ($Index_j$) is calculated as follows:

$$Index_j = 3(WI_j) - WF_j + WC_j$$

where WI_j measures the relative importance of road project *j* in the side constraints (the larger the value the less desirable the index), WF_j measures the relative amount of traffic flow on link *ab* for which project *j* is an option (the more the traffic the better the index), and WC_j measures the relative cost of road project *j* in the objective function (the larger the cost the less desirable the index). The computations for each component are described in detail below.

The WI_j Component—WI_j is the relative importance of candidate project *j* in the side constraints, weighted by 1 minus the activity level for road project *j* in the current LP solution. It is the sum of the relative impacts (IMP_{ij}) of project *j* in each of the side constraints:

$$WI_j = \sum_{i=1}^{n} IMP_{ij}$$

For constraints having only upper bounds

$$IMP_{ij} = [(1 - A_j)a_{ij}]/H_i$$

For constraints having only lower bounds

$$IMP_{ij} = [(1 - A_j)(-1)a_{ij}]/L_i$$

For constraints having both upper and lower bounds

$$IMP_{ij} = \text{the larger of: } [(1 - A_j)a_{ij}]/H_i \text{ or } [(1 - A_j)(-1)a_{ij}]/L_i$$

For equality constraints

$$IMP_{ij} = |(1 - A_j)a_{ij}| / E_i$$

where

a_{ij} = the coefficient for road project *j* in side constraint *i*

A_j = the activity level (solution value) for road project *j* in the current LP solution

H_i = the space available in side constraint *i*, which is a "less than" constraint. It is the right-hand-side value (upper bound) of constraint *i*, minus the activity level for constraint *i* (including the deviation variables), plus the upper bound for the lower-cost deviation variable for constraint *i*, plus 0.001. The last term is added to avoid the possibility of dividing by 0.

L_i = the space available in side constraint *i*, which is a "greater than" constraint. It is the activity level for constraint *i*, minus the right-hand-side value (lower bound) of constraint *i*, plus the upper bound for the lower-cost deviation variable for constraint *i*, plus 0.001.

E_i = a proxy for room remaining in side constraint *i*, which is an equality constraint. It is the right-hand-side value, plus the upper bound for the lower-cost deviation variables (these bounds are the same number) multiplied by 0.15. This

calculation attempts to produce a proxy having the same magnitude relative to its right-hand-side value as L_i and H_i have to their respective bounds. E_i is needed because an equality constraint will always be exactly satisfied (if the LP solution was feasible) and therefore will never have space remaining as do upper and lower bound constraints.

n = the number of side constraints.

The larger the value for WI_j, the less desirable is the index. The contribution of a_{ij} to the index is inversely proportional to activity level of project j in the current LP solution as a result of the $1 - A_j$ weights. Thus, the important road projects in the continuous LP (as measured by the magnitude of the activity level) have less penalty for side-constraint impact than do road projects having lower fractional values.

The WF_j Component—This component measures the amount of traffic flow observed on link ab relative to the average traffic flow on the network in the current LP solution. The larger this component, the lower (more desirable) the index. It is computed as:

$$WF_j = [\sum_{t=1}^{T} T_{ab_j,t}]/[(\sum_{ab\varepsilon L} \sum_{t=1}^{T} T_{ab,t})/M]$$

where

$T_{ab_j,t}$ = total amount of traffic flow (expressed in capacity units) on link ab_j (the link associated with road project j) in period t

$T_{ab,t}$ = total amount of traffic flow (expressed in capacity units) on link ab in period t

M = total number of nonzero values for $T_{ab,t}$ across all links on the network

T = number of periods

L = number of links.

The WC_j Component—This component is the relative cost of the road project, weighted by 1 minus the activity level for road project j in the current LP solution. It is computed as follows for maximization problems:

$$WC_j = [(1 - A_j)(-1)C_j] / [\sum_{j=1}^{J} |C_j| / N]$$

and for minimization problems:

$$WC_j = [(1 - A_j)C_j] / [\sum_{j=1}^{J} |C_j| / N]$$

where

C_j = the objective function coefficient for project j

A_j = the activity level for road project j in the current LP solution

N = the total number of nonzero C_j values

J = the total number of candidate road projects.

The greater this component, the larger (less desirable) the index. The weights $(1 - A_j)$ make the contribution of C_j to the index inversely proportional to the activity level for project j in the current LP solution. Thus, the more important projects (as measured by the size of their fractional value) have less penalty for construction cost.

"Tree-Growth" Rounding Procedure

The "tree-growth" rounding procedure derives its name from the order in which road projects are rounded. Only road projects that connect either directly to a demand node, or indirectly through previously rounded projects can be rounded to 1. As more road projects are rounded to 1, the growth of the "developed portion" of the network resembles tree growth. Unlike most trees, however, the branches originating from one demand node can connect (grow together) with branches from other demand nodes. This method of selecting road projects ensures that the "developed portion" of the network is comprised of one or more continuous networks, each of which connect to at least one final demand node. The following steps describe the rounding process that is applied in each iteration, beginning with iteration 3.

1. Initialize the "developed portion" of the road network—The "developed portion" of the network consists of road projects having integer values which are available for carrying traffic to the final demand nodes. At the beginning of an iteration, the "developed portion" of the network is comprised of three parts. First, it includes the road projects rounded to 1 in the previous iterations (in iteration 3, this consists only of the demand nodes). Second, the candidate road projects representing existing roads are rounded to 1 and added to the "developed portion" of the network. Third, any preselected road projects (user set them to 1) are added to the "developed portion" of the network.

2. Build the list of "connecting projects"—The "connecting projects" are the candidate projects that connect both spatially and temporally with the "developed portion" of the network. A candidate project connects spatially if it shares a node with one or more links included in the "developed portion" of the network. That project connects temporally if its period of implementation is greater than or equal to the period of implementation for at least one road construction project in the "developed portion" of the network with which the common node is shared. In other words, the road network must be developed to the common node by the period of implementation of a candidate project for that project to be feasible temporally.

3. Check the feasibility of the "connecting project" with the lowest index value—First identify the "connecting project" having the lowest (best) index value. If there are no remaining "connecting projects," proceed to step 8. Let X_{cur} represent the "connecting project" having the lowest (best) index value. The feasibility of rounding X_{cur} to 1 is checked for each side constraint. This is accomplished for constraint *i* by comparing the impact of X_{cur} plus the cumulative impacts from previous rounding decisions made in the current iteration against the "allotted space" for constraint *i*. "Allotted space" is the amount of impact from rounding decisions that is permitted in the current iteration. "Allotted space" intentionally underestimates the estimated "actual space" (the estimated actual amount of space available for absorbing impacts from rounded road projects) to avoid making too many decisions in any one iteration. (See appendix D for details regarding this feasibility check.)

(a) If the cumulative impact from rounding X_{cur} is *within* the "allotted space" *for all* side constraints, go to step 5.
(b) If the cumulative impact from rounding X_{cur} *exceeds* the "allotted space" *for any* side constraint, go to step 4.

4. Check the feasibility of X_{cur} against the "actual space" available in each side constraint—In this step the impacts of X_{cur} plus the cumulative impacts from the previous rounding decisions made in the current iteration are compared against the "actual space" available in each side constraint.

(a) If the cumulative impact associated with rounding X_{cur} is less than the "actual space" for all side constraints, then:

(i) if an alternate rounding choice (X_{alt}) has not yet been selected since last updating the list of connecting projects, let $X_{alt} = X_{cur}$,

(ii) otherwise, maintain the previously selected X_{alt}. (X_{alt} remains a candidate for rounding, and will eventually be rounded providing its index value proves to be substantially better than the next best feasible project.) Return to step 3 to process the "connecting project" with the next best index.

(b) If the cumulative impact associated with rounding X_{cur} exceeds the "actual space" for any one side constraint, remove X_{cur} from the "connecting project" list and add it to the "infeasible connecting project" list, then return to step 3 to process the next "connecting project."

5. Make a rounding decision—This step is reached if the impacts from X_{cur} were found by step 3 to be feasible for all side constraints (impacts are within "allotted space"). The index value for X_{cur}, ($INDEX_{cur}$) is compared with the index value for the alternate project ($INDEX_{alt}$) to determine which project should be rounded to 1. If the difference in index value is less than or equal to 0.5, X_{cur} is selected. In this case X_{alt} is not sufficiently better to warrant exceeding the threshold of acceptable impact for at least one side constraint. If the difference in index value exceeds 0.5, X_{alt} is selected over X_{cur} even though it will have a significant impact on at least one side constraint. The details of this step follow:

(a) *If:* $INDEX_{cur} - INDEX_{alt} \leq 0.5$, then round X_{cur} to 1. Next, go to step 6.
(b) *If:* $INDEX_{cur} - INDEX_{alt} > 0.5$, then round X_{alt} to 1. Next, go to step 10 (no further rounding decisions will be made in the current iteration because X_{alt} exceeds the "allotted space" for at least one side constraint).

6. Update various lists—After a road project has been rounded to 1, various lists must be updated. First, the rounded project is added to the "developed portion" of the network. Second, the impacts of the rounded project are added to the cumulative totals for the side constraints. Third, the list of "connecting projects" is updated as follows:

(a) The rounded project is removed.
(b) The candidate projects associated with fractional links connecting to the rounded project are added.
(c) The "infeasible connecting projects" and X_{alt} are added. There is a possibility that the rounded project created more "actual space" (rather than less) in critical side constraints. So, the previously infeasible projects are added back to "connecting project" list because they may now be feasible.

Proceed to step 7.

7. Has the maximum number of rounding decisions for the current iteration been reached?—The maximum number of projects that can be rounded in iterations 3 and 4 equals 50 percent of the original candidate projects plus 1. This is increased to 75 percent plus 1 beginning with iteration 5. These restrictions represent a compromise. Experience suggests that better solutions are obtained when a small number of rounding decisions are made in each iteration. This permits the ensuing LP solution to react to these rounding decisions prior to making additional decisions. On the other hand, making as many decisions as possible in each iteration reduces the number of iterations, thereby reducing the amount of computer resources required for a HIP solution.

(a) If the maximum number of rounding decisions has not been reached, go back to step 3.
(b) If the maximum number of rounding decisions has been reached, proceed to step 10.

8. Round the alternative project (X_{alt}), if one has been identified since updating the list of "connecting projects"—This step is reached (from step 3) if none of the remaining "connecting projects" is within the "allotted space," but space for absorbing impacts remains in the side constraints and the maximum number of rounding decisions for the current iteration has not yet been reached.

(a) If there is an alternative project, X_{alt}, round this project to 1. Then proceed to step 10.
(b) If there is no X_{alt} available (the impact of each connecting project plus the cumulative impact of all previously rounded projects exceeds the "actual space" in each of the side constraints), proceed to step 9.

9. Have any rounding decisions been made in the current iteration?—This includes modifications made in previous rounding decisions (road standard changes, etc.) as well as whether any projects have been rounded to 1 in the current iteration.

(a) If at least one rounding decision has been made in the current iteration, proceed to step 10.
(b) If no rounding decisions have been made in the current iteration, then the impacts from each of the "connecting projects" exceed the "actual space" in each of the side constraints. Yet, fractional projects remain. Something must be done to avoid getting the same result in the next iteration. Implement the following rules and then proceed to step 10.
   (i) If either the objective function value or the list of candidate projects differ from the previous iteration, then each of the "connecting projects" are set to 0. In addition, the road projects of the same road standard applying to earlier periods on those links are also set to 0.
   (ii) If both the objective function value and list of candidate projects are the same as the previous iteration, then set to 0 all of the candidate projects (whether they connect or not), plus the road projects of the same road standard applying to earlier periods on those links.

10. End the rounding procedure in the current iteration.

APPENDIX D: FEASIBILITY CHECKING

In the HIP process the feasibility of rounding option j in side constraint i is checked against either "actual space" or "allotted space." These checks are made in the rounding procedure (see the discussion for iteration 3 in the main report, and/or appendix C) and in modifications made to previous rounding decisions (see the discussion for iterations 4, 5, and 6 in the main report).

Comparing Against the "Actual Space"

"Actual space" measures the total amount of impact that constraint i can absorb from rounding decisions in the current iteration. The feasibility of rounding option j with regard to the "actual space" in side constraint i is checked by comparing its impact ($IMPACT_{ij}$) plus the cumulative impacts on constraint i of previous rounding decisions in the current iteration (CUM_i), against the "actual space" in constraint i ($SPACE_i$).

$IMPACT_{ij}$ is calculated for the various types of rounding options as follows:

1. The impact on constraint i of rounding road project j to 1 is:

$$IMPACT_{ij} = a_{ij}$$

where a_{ij} is the coefficient of project j in side constraint i.

2. The impact on constraint i of a change in road standard or postponing construction in the "developed portion" of the network is:

$$IMPACT_{ij} = (a_{ij(\text{new})}) - (a_{ij(\text{old})})$$

where $a_{ij(\text{new})}$ is the coefficient in constraint i for the road project replacing the previously selected project, whose coefficient in constraint i is $a_{ij(\text{old})}$.

3. The impact on constraint i of closing a link (by setting the previously rounded road project to 0) is:

$$IMPACT_{ij} = -(a_{ij(\text{old})})$$

where $a_{ij(\text{old})}$ is the coefficient in constraint i of the previously selected project.

CUM_i is the cumulative total impact on constraint i of the rounding decisions made in the current iteration. The impact of rounding option j is added to CUM_i when rounding option j is selected and incorporated in the LP matrix.

Feasibility for Upper Bound Constraints—The available space in side constraint i for absorbing impacts from heuristic rounding decisions, $SPACE_i$, is calculated at the beginning of a heuristic iteration (immediately following a LP solution). For upper bound constraints the calculation is:

$$SPACE_i = RHS_i - [A_i + D_{iL}^{U} + D_{iH}^{U} - \sum_j a_{ij}(R_j)] \qquad \text{(D-1)}$$

where

RHS_i	=	the right-hand-side value for side constraint i
A_i	=	the activity level for side constraint i (includes the effects of the deviation variables)
D_{iL}^{U}	=	the activity level (solution value) for the lower-cost deviation variable associated with the upper bound placed on constraint i
D_{iH}^{U}	=	the activity level (solution value) for the higher-cost deviation variable associated with the upper bound placed on constraint i

a_{ij} = the coefficient for road project j in side constraint i

R_j = the activity level (solution value) for road project j, which is in the current LP solution at some fractional value.

For upper bound constraints, a positive value for $SPACE_i$ indicates there is "actual space" available for absorbing impacts from rounding decisions, a negative value indicates that the upper bound on constraint i was violated in the current LP solution (there is, in fact, negative "actual space"), and a value of zero indicates the constraint was at its bound (no space is available).

Rounding options with negative $IMPACT_{ij}$ values decrease the cumulative total impact from rounding decisions (CUM_i). These options are, therefore, always feasible, regardless of the value of $SPACE_i$ or CUM_i.

Rounding options with positive values for $IMPACT_{ij}$ increase the cumulative total impact from rounding decisions (CUM_i). These options are feasible if

$$IMPACT_{ij} + CUM_i \leq SPACE_i$$

The cumulative total impact of rounding decisions on side constraint i, CUM_i, must be updated following each rounding decision. This is done by adding $IMPACT_{ij}$ for the selected rounding option to CUM_i.

Feasibility for Lower Bound Constraints—The available space with regard to a lower bound placed on side constraint i is calculated as follows:

$$SPACE_i = RHS_i - [A_i - D_{iL}^{L} - D_{iH}^{L} - \sum_j a_{ij}(R_j)] \tag{D-2}$$

where D_{iL}^{L} is the activity level for the lower-cost deviation variable associated with the lower bound on constraint i, D_{iH}^{L} is the activity level for the higher-cost deviation variable associated with the lower bound on constraint i, and all other terms are as previously defined.

A negative value for $SPACE_i$ indicates there is "actual space" available for absorbing impacts from rounding decisions, a positive value indicates that the lower bound on constraint i was violated in the current LP solution (there is, in fact, negative "actual space"), and a value of 0 indicates constraint i was at its lower bound (no space is available).

Impacts for rounding options are calculated the same way as for upper-bound constraints, but the test for identifying feasible rounding options is reversed. Rounding options with positive $IMPACT_{ij}$ values increase CUM_i, thereby moving it away from $SPACE_i$. These rounding options are, therefore, always feasible for lower bound constraints, regardless of the values of $SPACE_i$ or CUM_i.

Rounding options with negative values for $IMPACT_{ij}$ decrease CUM_i thereby moving it closer to the "actual space" for constraint i. Rounding options with a negative $IMPACT_{ij}$ are feasible if

$$IMPACT_{ij} + CUM_i \geq SPACE_i$$

The cumulative total impact of rounding decisions on side constraint i, CUM_i, must be updated following each rounding decision. This is done by adding $IMPACT_{ij}$ for the selected rounding option to CUM_i.

Feasibility for Equality Constraints—Equality constraints have upper and lower bounds set to the same right-hand-side value. Feasibility with regard to both bounds is checked. A rounding decision is considered infeasible if the calculations for either the upper or lower bounds indicate an infeasibility.

Feasibility for Constraints Having Ranges—Constraints for which "ranges" have been specified contain both lower and upper bounds. Feasibility with regard to both bounds is checked. A rounding decision is considered infeasible if the calculations for either the upper or lower bounds indicate an infeasibility.

Comparing Against the "Allotted Space"

The "allotted space" is computed by adjusting "actual space." When road construction plays a significant role in side constraint i (10 percent or more of the impacts in the constraint are associated with road construction), "actual space" is adjusted downward to compute "allotted space." In these instances "allotted space" is a more conservative assessment of the amount of rounding-decision impacts to permit on constraint i in the current iteration. Experience indicates that consuming the available space in a side constraint gradually over a moderate number of iterations is better than using it all in one iteration. This gives the LP solver an opportunity to react to the rounding decisions, and to provide an improved basis for making additional rounding decisions.

When road construction comprises less than 10 percent of the side constraint impact, "actual space" is adjusted upward a small amount to arrive at "allotted space." In these cases, roads play a sufficiently minor role that impacts exceeding "actual space" can be accommodated in subsequent LP solutions by adjustments in the resource projects and traffic flow variables selected.

Feasibility for Upper Bound Constraints—The "allotted space" for an upper bound on constraint i ($ALLOT_{UB,i}$) is calculated at the beginning of an iteration as follows:

$$ALLOT_{UB,i} = SPACE_i + SPACE_i\ [(DELTA_i)(ADJUST)]$$

where $SPACE_i$ is as defined earlier. $DELTA_i$ and $ADJUST$ are adjustment factors whose computations are described later in this appendix.

A positive value for $ALLOT_{UB,i}$ indicates there is space available in side constraint i. Rounding options with positive values for $IMPACT_{ij}$ increase the cumulative total impact from rounding decisions (CUM_i). Rounding option j having a positive $IMPACT_{ij}$ is within the "allotted space" if

$$IMPACT_{ij} + CUM_i \leq ALLOT_{UB,i}$$

where all terms are as previously defined.

Rounding options with negative $IMPACT_{ij}$ values decrease the cumulative total impact from rounding decisions (CUM_i). These options are, therefore, always feasible, regardless of the value of $ALLOT_{UB,i}$ or CUM_i.

Feasibility for Lower Bound Constraints—The "allotted space" for a lower bound on constraint i ($ALLOT_{LB,i}$) is calculated at the beginning of an iteration as follows:

$$ALLOT_{LB,i} = SPACE_i\ - SPACE_i\ [(DELTA_i)(ADJUST)]$$

where $SPACE_i$ is as defined earlier, and $DELTA_i$ and $ADJUST$ are adjustment factors whose computations are described later in this appendix.

A negative value for $ALLOT_{LB,i}$ indicates there is space available in side constraint i. Rounding options with negative values for $IMPACT_{ij}$ decrease CUM_i, thereby moving it closer to the "actual space" for constraint i. Rounding options with a negative $IMPACT_{ij}$ are feasible if

$$IMPACT_{ij} + CUM_i \geq ALLOT_{LB,i}$$

where all terms are as previously defined.

Rounding options with positive $IMPACT_{ij}$ values increase the cumulative total impact CUM_i. They are, therefore, always feasible, regardless of the value of $ALLOT_{LB,i}$ or CUM_i.

Threshold Feasibility for Equality Constraints—Equality constraints have upper and lower bounds set to the same right-hand-side value. Feasibility with regard to both bounds is checked. A rounding decision is considered infeasible if the calculations for either the upper or lower bounds indicate an infeasibility.

Feasibility for Constraints Having Ranges—Constraints for which "ranges" have been specified contain both lower and upper bounds. Feasibility with regard to both bounds is checked. A rounding decision is considered infeasible if the calculations for either the upper or lower bounds indicate an infeasibility.

$DELTA_i$—This term is an adjustment factor based on the relative importance of road construction in side constraint i. In general, this factor makes the threshold for constraint i more restrictive the greater the importance of road construction in constraint i. $DELTA_i$ is calculated as:

$$DELTA_i = \text{smaller of } \{0.1, \text{ or } -0.5 + 0.05(1/BETA_i)\}$$

$BETA_i$ is a measure of the relative importance of road construction in side constraint i, and is calculated as:

$$BETA_i = \sum_{j\varepsilon S} |a_{ij}| \,/\, [\sum_{k} |a_{ik} X_k| + \sum_{m} |a_{im} F_m| + \sum_{j\varepsilon S} |a_{ij}|]$$

where

a_{ij} = the coefficient for road project j in constraint i

S = the set of road projects that have a nonzero activity level in the current LP solution

a_{ik} = the coefficient for resource project k in constraint i

X_k = the activity level for resource project k in the current LP solution

a_{im} = the coefficient for traffic variable m in constraint i

F_m = the activity level for traffic variable m in the current LP solution.

The range of $DELTA_i$ with respect to $BETA_i$ is shown below:

$DELTA_i$	$BETA_i$
–0.45	1.000
–.43	.750
–.40	.500
–.30	.250
–.17	.150
0	.100
.10	≤.083

$BETA_i$, the relative impact of road construction on constraint i, is inversely related to $DELTA_i$. This means "allotted space" is a smaller percentage of "actual space" ($SPACE_i$) the larger the relative impact of road construction. This avoids making too many rounding decisions in one iteration. Experience indicates that better final solutions are obtained if the number of rounding decisions made in any one iteration are somewhat limited. Limiting based on constraint i, however, becomes less important for smaller

values for $BETA_i$. When the importance of road construction in constraint i is less, the LP solution process can more easily adjust to the impacts of rounding decisions by its selection of resource projects and traffic routing in later iterations. In these instances, "allotted space" is less restrictive to avoid unnecessary iterations in the heuristic process.

ADJUST—Tests indicated that the limits imposed by $DELTA_i$ were too restrictive in the later iterations. Too many iterations were required to achieve integer values for the road projects. The factor *ADJUST* was added to alleviate this problem. This factor reduces the effect of $DELTA_i$ as the number of iterations increase. *ADJUST* is calculated as follows:

$$ADJUST = \text{larger of } \{0, \text{ or } 1.2 - 0.2\,(IT\text{–}2)\}$$

where *IT* is the number of iterations. *IT*–2 is the number of iterations in which rounding decisions have been made. (Iterations 1 and 2 are not counted because they do not involve rounding decisions.) *ADJUST* equals zero beginning with iteration 8.

☆U.S. GOVERNMENT PRINTING OFFICE:1991 -573-041/41008

Jones, J. Greg; Weintraub, Andres; Meacham, Mary L.; Magendzo, Adrian. 1991. A heuristic process for solving mixed-integer land management and transportation planning models. Res. Pap. INT-447. Ogden, UT: U.S. Department of Agriculture, Forest Service, Intermountain Research Station. 39 p.

Authors describe a heuristic procedure for solving a linear mixed-integer mathematical programming formulation useful for analyzing where and when to conduct specific land management activities and road construction and reconstruction projects. The procedure develops feasible mixed-integer solutions with objective values generally within 10 percent of the mixed-integer optimal solutions, but requires substantially less computer time than the standard branch-and-bound approach for solving this mixed-integer formulation.

KEYWORDS: heuristic decision rules, heuristic integer programming, mixed-integer programming, land management planning, transportation planning, network analysis

Printed on recycled paper

INTERMOUNTAIN RESEARCH STATION

The Intermountain Research Station provides scientific knowledge and technology to improve management, protection, and use of the forests and rangelands of the Intermountain West. Research is designed to meet the needs of National Forest managers, Federal and State agencies, industry, academic institutions, public and private organizations, and individuals. Results of research are made available through publications, symposia, workshops, training sessions, and personal contacts.

The Intermountain Research Station territory includes Montana, Idaho, Utah, Nevada, and western Wyoming. Eighty-five percent of the lands in the Station area, about 231 million acres, are classified as forest or rangeland. They include grasslands, deserts, shrublands, alpine areas, and forests. They provide fiber for forest industries, minerals and fossil fuels for energy and industrial development, water for domestic and industrial consumption, forage for livestock and wildlife, and recreation opportunities for millions of visitors.

Several Station units conduct research in additional western States, or have missions that are national or international in scope.

Station laboratories are located in:

Boise, Idaho

Bozeman, Montana (in cooperation with Montana State University)

Logan, Utah (in cooperation with Utah State University)

Missoula, Montana (in cooperation with the University of Montana)

Moscow, Idaho (in cooperation with the University of Idaho)

Ogden, Utah

Provo, Utah (in cooperation with Brigham Young University)

Reno, Nevada (in cooperation with the University of Nevada)